AMD RADEON RX 590
The affordable GPU for today's gaming **PG. 74**

SOLID-STATE NATION
Find the perfect SSD for your system & budget **PG. 36**

ACER PREDATOR X27
Is this the best screen you can buy right now? **PG. 82**

MAXIMUMPC

MINIMUM BS • JANUARY 2019 • www.maximumpc.com

2019 TECH PREVIEW

Next-gen hardware from AMD, Intel, & Nvidia that will transform your PC

PLUS!
Build a **1080p gaming rig** **PG. 68**

THE BEST GAMES COMING THIS YEAR

Revealed: The biggest & finest titles for 2019 **PG. 48**

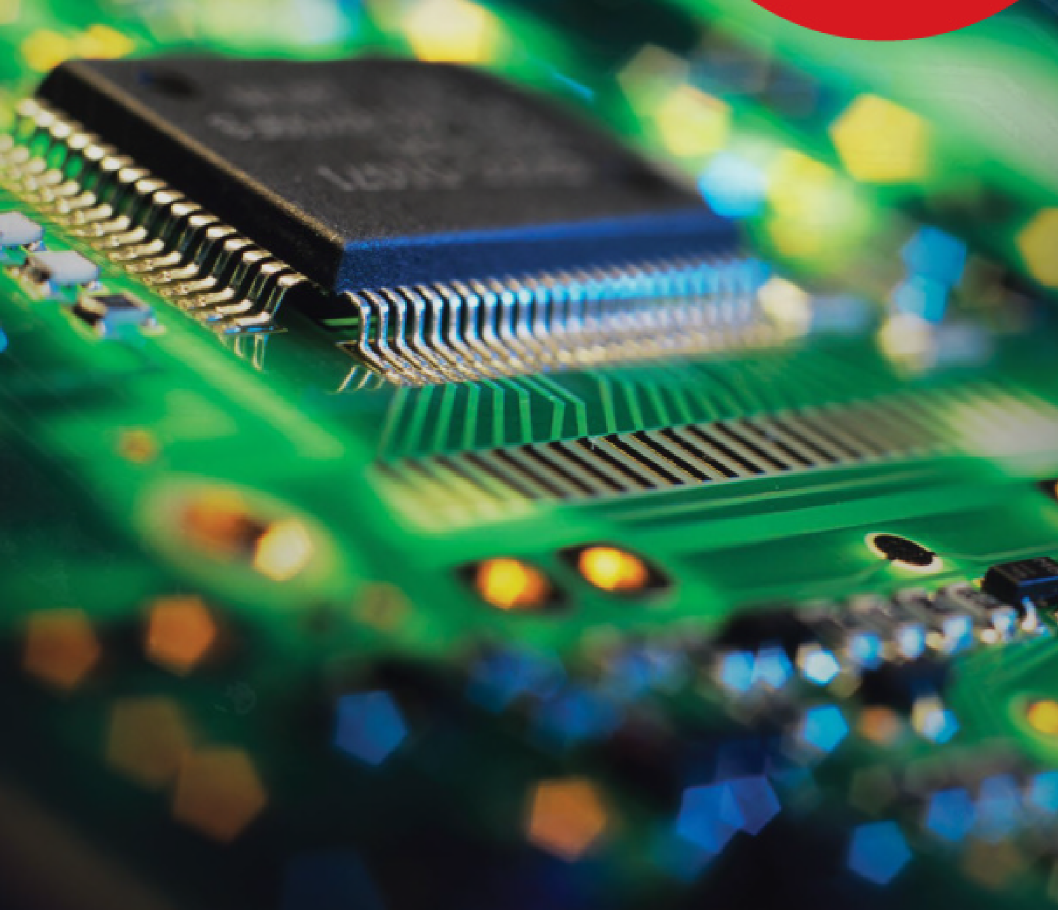

FUTURE

CORSAIR
TYPE ALL DAY
PLAY ALL NIGHT

/CORSAIR

K70 RGB MK.2 LOW PROFILE

RAPIDFIRE Mechanical Gaming Keyboard

Experience all-day comfort without sacrificing mechanical gaming performance. The CORSAIR K70 RGB MK.2 LOW PROFILE's CHERRY MX Low Profile RGB Speed mechanical keyswitches provide the comfort, reliability, performance and accuracy you demand with ultra-fast 1.0mm actuation.

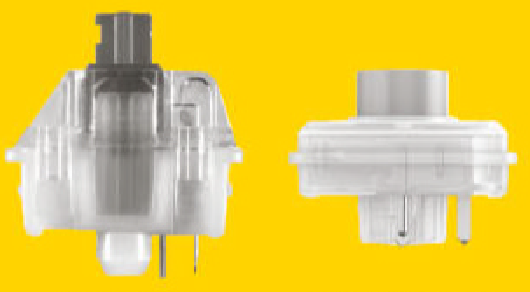

STANDARD
MX SPEED SWITCH

MX LOW PROFILE
SPEED SWITCH

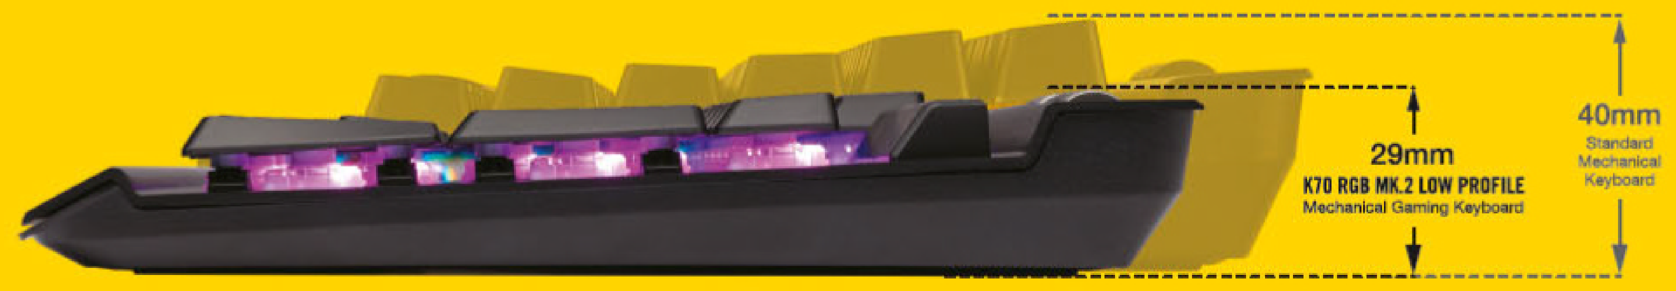
29mm
K70 RGB MK.2 LOW PROFILE
Mechanical Gaming Keyboard
40mm
Standard
Mechanical
Keyboard

Experience the best lighting effects with CORSAIR iCUE Software
Windows® 7, 8.1, and 10 compatible

FIND OUT MORE AT CORSAIR.COM

PC GAMER
"Put together by genuine enthusiasts with mastery of the craft."
NEW
NEURON
GEFORCE RTX

PC GAMER EDITOR'S CHOICE ★★★★★

MAXIMUM PC KICK ASS!

intel CORE i9 9th Gen

Powered by Intel®
Core™ i9 Processors

FREE Lifetime 24/7 Service & Support
Online Financing Available

Intel, the Intel logo, Intel Inside, Intel Core are trademarks of Intel Corporation or its subsidiaries in the U.S. and/or other countries.

© 2018 NVIDIA Corporation, NVIDIA, the NVIDIA logo, GeForce, and NVIDIA Turing are trademarks or registered trademarks of NVIDIA Corporation.
All rights reserved.

The ORIGIN of Performance
EVO17-S
EON15-S
ORIGIN

ORIGINPC.COM

@originpc

ORIGIN
HIGH-PERFORMANCE PCS

© 2018 ORIGIN PC Corporation. All rights reserved. ORIGIN "O" Symbols are trademarks of ORIGIN PC Corporation.
All other trademarks are property of their respective owners.

SUBSCRIBE TODAY!
see PG. 44

inside

JANUARY 2019

24 TECH PREVIEW 2019
Next-gen hardware that will transform your PC

32 QUANTUM COMPUTING
An introduction to the coolest tech around; learn how it works, and how to get started.

36 SOLID STATE OF THE NATION
With SSD prices at an all-time low, we evaluate your solid-state storage options.

48 GAMES FOR 2019
From heavy gunfire to forest spirits, here are our top tips for the best games of 2019.

QUICKSTART

12 THE NEWS
AMD ups the pressure on Intel; Intel's new top dog; Facebook to face the music; Microsoft tests ads; more....

18 THE LIST
The top eight free web apps that are just as good as desktop ones.

What happens when AMD accidentally sends you two RX 590s...

R&D

56 AUTOPSY
We take a microscopic look at the Mac mini 2018's innards.

58 HOW TO
Program and set up a tiny web server; build a Pi-powered random video player; fix OneDrive sync problems; create an animated GIF.

68 BUILD IT
How we pieced together our Full HD AMD gaming rig.

LETTERS

22 DOCTOR

94 COMMENTS

IN THE LAB

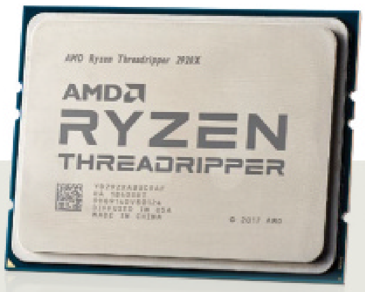

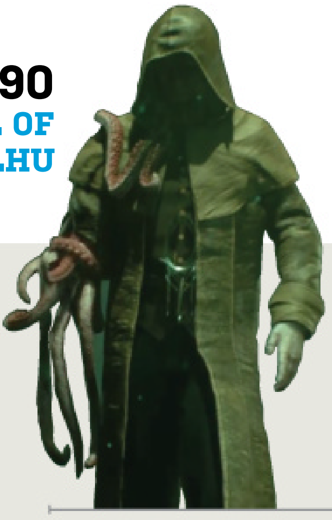

76 AMD THREADRIPPER 2920X

80 ASUS ROG STRIX GL12CX COD EDITION

85 CRUCIAL BX500 120GB

89 THERMALTAKE TT PREMIUM X1 RGB

90 CALL OF CTHULHU

ENERMAX
POWER. INNOVATION. DESIGN.

RGB Addressable ENERMAX · ASUS AURA SYNC · ASRock POLYCHROME SYNC · GIGABYTE RGB FUSION READY · RGB msi MYSTIC LIGHT SYNC

SABERAY ADV

SABERAY White

- Support Addressable RGB Lighting Sync
- 3x ENERMAX Unique Square Fans Pre-installed
- Acrylic and Meshed Front Panels Included
- Free Gift - Monitor Stand included

@Enermaxusa | www.Enermaxusa.com

Available Now at Fry's Electronics amazon newegg.com

SAGER
GAMING NOTEBOOKS

SAGER RECOMMENDS WINDOWS

BLAZING PERFORMANCE

NP9877 NOTEBOOK

$2,999

FREE UPS GROUND SHIPPING*

- **Windows 10 Home**
- **9th Generation Intel® Core™ i7-9700K Processor**
- Opt. 9th Generation Intel® Core™ i9-9900K Processor
- 17.3" Full HD, 144Hz Matte Display with G-SYNC
- Opt. 17.3" 4K QFHD Matte Display with G-SYNC
- NVIDIA® GeForce™ GTX 1080 8GB GPU
- Opt. Dual NVIDIA® GeForce™ GTX 1080 GPUs
- 16GB DDR4-2400MHz Memory
- 250GB WD Black NVMe M.2 SSD + 1TB 7200RPM HDD
- Killer Wireless-AC 1550 + Bluetooth

$2,399

NP9155 NOTEBOOK

FREE UPS GROUND SHIPPING*

- **Windows 10 Home**
- **9th Generation Intel® Core™ i7-9700K Processor**
- Opt. 9th Generation Intel® Core™ i9-9900K Processor
- 15.6" Full HD 144Hz Matte Display with G-SYNC
- Opt. 17.3" FHD 144Hz or 4K QFHD Matte Display (NP9175)
- NVIDIA® GeForce™ GTX 1070 8GB GPU
- Opt. NVIDIA® GeForce™ GTX 1060 or 1080 GPU
- 16GB DDR4-2400MHz Memory
- 250GB WD Blue M.2 SSD + 1TB 7200RPM HDD
- Killer Wireless-AC 1550 + Bluetooth

$1,899

NP8955 NOTEBOOK

FREE UPS GROUND SHIPPING*

- **Windows 10 Home**
- **8th Generation Intel® Core™ i7-8750H Processor**
- 15.6" Full HD 144Hz 5ms IPS Matte Display
- NVIDIA® GeForce™ GTX 1070 8GB GPU with Max-Q
- Opt. GeForce™ GTX 1060 6GB GPU (NP8953)
- 16GB DDR4-2666MHz Memory
- 250GB WD Black NVMe M.2 SSD + 1TB 7200RPM HDD
- Full Color Programmable Backlight Keyboard
- Intel® Dual Band Wireless-AC 9560 M.2 AC + BT
- Ultra-slim design with only 0.73 inches thin

All Sager systems come with one year Limited Warranty & Lifetime Tech Support. Prices and specifications are subject to change without notice or obligation. Sager Corporate Offices 18005 Cortney Court, City of Industry, California 91748 ©2018 by Midern Computer, Inc. All rights reserved. Intel, the Intel Logo, Intel Inside, Intel Core, and Core Inside are trademarks of Intel Corporation or its subsidiaries in the U.S. and/or other countries. All company and/or product names mentioned herein are trademarks and/or registered trademarks of their respective companies. *Free UPS Ground Shipping valid to contiguous US order only.

Windows 10

intel
CORE i7
8th Gen

$1,999

NP8375 NOTEBOOK FREE UPS GROUND SHIPPING*

- **Windows 10 Home**
- **8th Generation Intel® Core™ i7-8750H Processor**
- 17.3" Full HD 144Hz Matte Display with G-SYNC
- NVIDIA® GeForce™ GTX 1070 8GB GPU
- Opt. GeForce™ GTX 1060 6GB GPU (NP8373)
- 16GB DDR4-2666MHz Memory
- 250GB Samsung 970 EVO NVMe M.2 SSD + 1TB 7200RPM HDD
- Per-Key RGB Backlight Keyboard
- Intel® Dual Band Wireless-AC 9560 M.2 AC + BT
- Slim design with only 0.98 inches thin

$1,599

NP8954 NOTEBOOK After $200 Instant Savings!

- **Windows 10 Home**
- **8th Generation Intel® Core™ i7-8750H Processor**
- 15.6" Full HD 120Hz 3ms 94% NTSC Matte Display
- NVIDIA® GeForce™ GTX 1070 8GB GPU with Max-Q
- 16GB DDR4-2666MHz Memory
- 120GB WD Green M.2 SSD + 1TB 7200RPM HDD
- Full Color Programmable Backlight Keyboard
- Intel® Dual Band Wireless-AC 9560 M.2 AC +BT
- Ultra-slim design with only 0.73 inches thin

$1199

NP7851 NOTEBOOK After $100 Instant Savings!

- **Windows 10 Home**
- **8th Generation Intel® Core™ i7-8750H Processor**
- 15.6" Full HD 144Hz 5ms IPS Matte Display
- Opt. 17.3" Full HD 144Hz Matte Display (NP7871)
- NVIDIA® GeForce™ GTX 1060 6GB GPU
- 8GB DDR4-2666MHz Memory
- 250GB WD Blue M.2 SSD
- Full Color Programmable Backlight Keyboard
- Intel® Dual Band Wireless-AC 9560 M.2 AC + BT

$999

NP6853 NOTEBOOK After $100 Instant Savings!

- **Windows 10 Home**
- **8th Generation Intel® Core™ i7-8750H Processor**
- 15.6" Full HD 120Hz 3ms 94% NTSC Matte Display
- Opt. 17.3" FHD IPS Matte Display (NP6873)
- NVIDIA® GeForce™ GTX 1050Ti 4GB GPU
- 8GB DDR4-2666MHz Memory
- 1TB 5400RPM HDD
- Full Color Programmable Backlight Keyboard
- Intel® Dual Band Wireless-AC 9560 M.2 AC + BT

CALL US TOLL FREE: 1.800.669.1624 CUSTOMIZE & PRICING GO ONLINE: SAGERNOTEBOOK.COM

Dealer/Var, government and corporate pricing are available. Please call for details.

techradar.pro

IT INSIGHTS FOR BUSINESS

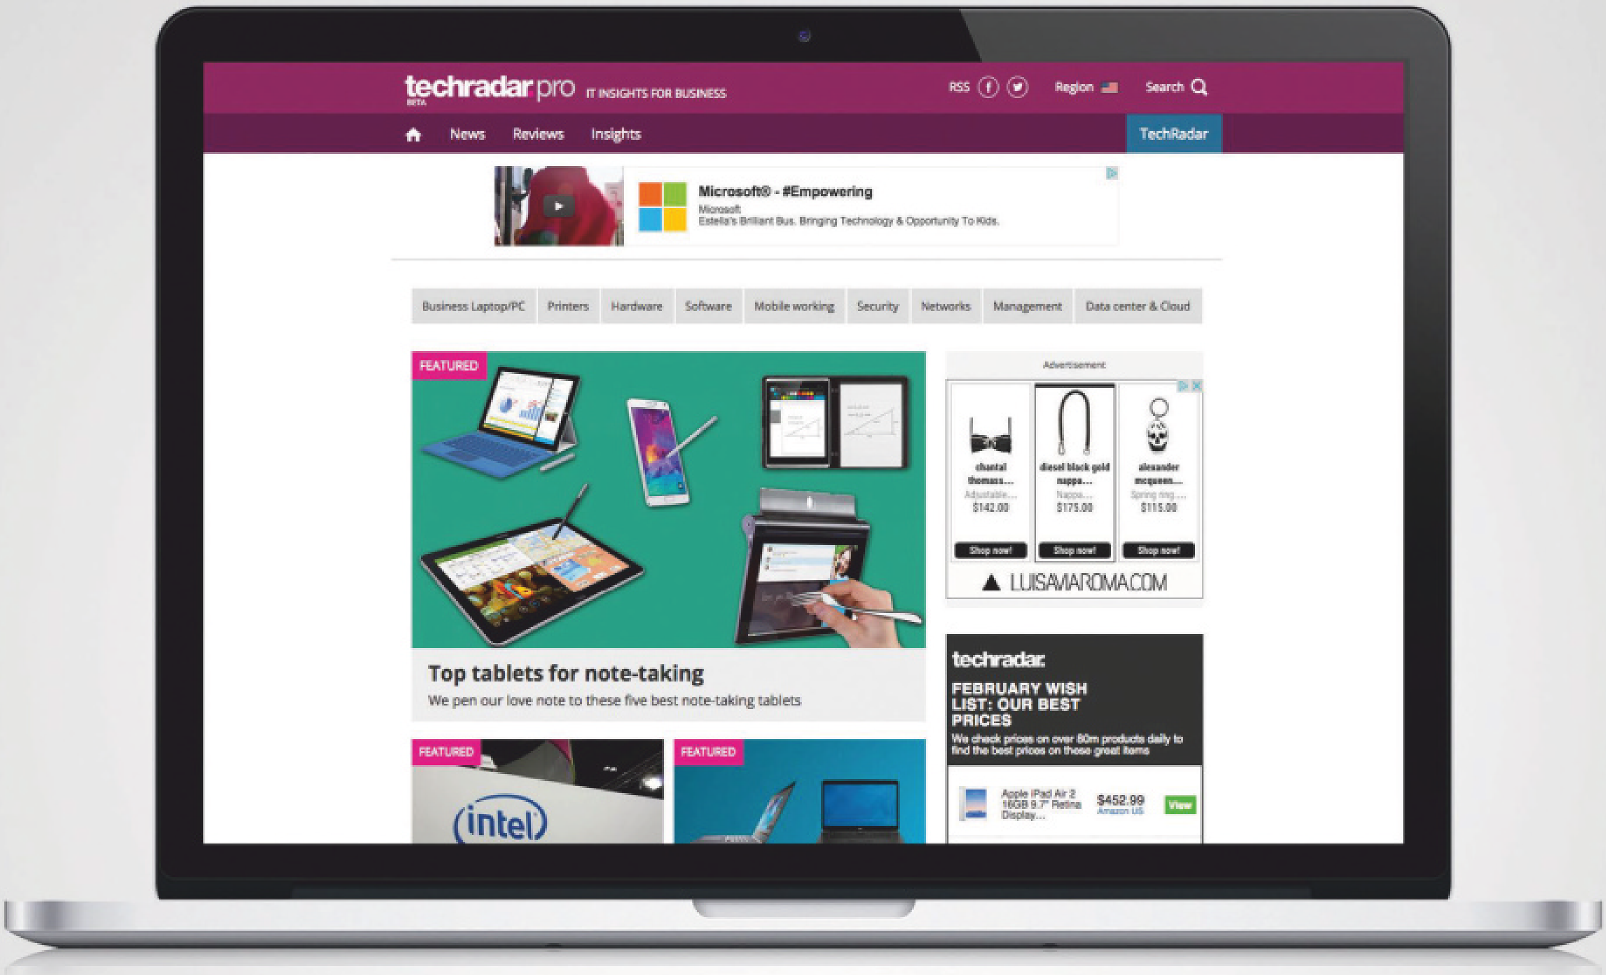
techradar pro
IT INSIGHTS FOR BUSINESS
RSS
Region
Search
News
Reviews
Insights
TechRadar
Microsoft® - #Empowering
Business Laptop/PC
Printers
Hardware
Software
Mobile working
Security
Networks
Management
Data center & Cloud
FEATURED
Advertisement
chantal thomass...
$142.00
diesel black gold nappa...
$175.00
alexander mcqueen...
$115.00
Shop now!
LUISAVIAROMA.COM
techradar.
Apple iPad Air 2 16GB 9.7" Retina Display...
$452.99
Amazon US
View
intel

Microsoft
Estella's Brilliant Bus. Bringing Technology & Opportunity To Kids.

Top tablets for note-taking

We pen our love note to these five best note-taking tablets

FEBRUARY WISH LIST: OUR BEST PRICES

We check prices on over 80m products daily to find the best prices on these great items

Best practices and solutions

Up-to-the minute news

In-depth reviews

http://www.techradar.com/pro

MAXIMUMPC

EDITORIAL
Executive Editor: Alan Dexter
Deputy Editor: Zak Storey
Senior Editor: Jarred Walton
Hardware Lead: Bo Moore
Hardware Staff Writer: Joanna Nelius
Contributing Editor: Chris Angelini
Contributing Writers: Mats Tage Axelsson, Alex Campbell, Alex Cox, Ian Evenden, Matt Hanson, Jeremy Laird, Chris Lloyd
Copy Editor: Katharine Davies
Editor Emeritus: Andrew Sanchez

ART
Art Editor: Fraser McDermott
Image Manipulation: Gary Stuckey
Photography: Future Photo Studio

BUSINESS
US Marketing & Strategic Partnerships: Stacy Gaines, stacy.gaines@futurenet.com
US Chief Revenue Officer: Luke Edson, luke.edson@futurenet.com
East Coast Account Director: Brandie Rushing, brandie.rushing@futurenet.com
East Coast Account Director: Michael Plump, michael.plump@futurenet.com
East Coast Account Director: Victoria Sanders, victoria.sanders@futurenet.com
East Coast Account Director: Melissa Planty, melissa.planty@futurenet.com
East Coast Account Director: Elizabeth Fleischman, elizabeth.fleischman@futurenet.com
West Coast Account Director: Austin Park, austin.park@futurenet.com
West Coast Account Director: Jack McAuliffe, jack.mcauliffe@futurenet.com
Director, Client Services: Tracy Lam, tracy.lam@futurenet.com

PRODUCTION
Head of Production: Mark Constance
Production Manager: Vivienne Calvert
Project Manager: Clare Scott
Production Assistant: Emily Wood

FUTURE US, INC.
11 West 42nd Street, 145th Floor, New York, NY 10036, USA
www.futureus.com

SUBSCRIBER CUSTOMER SERVICE
Maximum PC Customer Care,
Future Publishing, PO Box 5852, Harlan, IA 51593-1352
Website: http://myfavoritemagazines.com
Tel: 844-779-2822
Email: contact@myfavoritemagazines.com

BACK ISSUES
Website: http://myfavoritemagazines.com
Tel: +44 344 848 2852

Next Issue On Sale February 5, 2019

©2019 Future US, Inc. All rights reserved. No part of this magazine may be used or reproduced without the written permission of Future US, Inc. (owner). All information provided is, as far as Future (owner) is aware, based on information correct at the time of press. Readers are advised to contact manufacturers and retailers directly with regard to products/services referred to in this magazine. We welcome reader submissions, but cannot promise that they will be published or returned to you. By submitting materials to us, you agree to give Future the royalty-free, perpetual, non-exclusive right to publish and reuse your submission in any form, in any and all media, and to use your name and other information in connection with the submission.

FUTURE
Connectors.
Creators.
Experience
Makers.

Future plc is a public company quoted on the London Stock Exchange (symbol: FUTR)
www.futureplc.com

Chief executive **Zillah Byng-Thorne**
Non-executive chairman **Richard Huntingford**
Chief financial officer **Penny Ladkin-Brand**
Tel +44 (0)1225 442 244

Alan
Dexter

PREDICTING THE FUTURE ISN'T EASY, BUT WE LIKE A CHALLENGE

HUMANS LIKE to know what's coming. Look back through history, and it's littered with future-gazers, augurs, oracles, and prophets trying to guess what's going to happen next. We like to think we've come a long way from casting bones and imbibing noxious substances in order to foresee what the gods have in store for us, but realistically, a lot of our divinations have just gone a bit more high tech.

Supercomputers may be the stuff of tech dreams, but much of the time, they're trying to predict the future. From analyzing weather data to ascertain which way a storm is going, to trying to spot trends in the shares market, our greatest technological wonders are dedicated to trying to guess what's going to happen next. We have whole industries built on trying to foretell the future, and as it underpins the money markets, this means it affects pretty much every other industry, too.

When it comes to our field of interest, knowing what's around the corner can make a big difference to our purchasing and building decisions. If you're aware that a new processor is on the way, you can hold off that build you've been planning until the reviews drop. The same goes for graphics cards, memory, SSDs, screens—you name it, if you know something bigger and better is on the horizon, it's often worth waiting to see what sort of impact that will have before you start building.

We're in a slightly better position than most soothsayers, because the tech industry likes trends, and there's plenty of information around that shows exactly where things are going. We know, for example, that at some point we're going to move to DDR5, PCIe 4.0, USB 3.2, and numerous other advances. That's because a lot of the technologies we rely on are underpinned by well-documented standards. But what of the specifics? Turn to page 24 to see what our crystal balls have thrown up.

While we're in the mindset of trying to predict what's going to be the next big thing, this issue we also take a look at the games that are lining up to keep the PC at the top of the pile. Whether you're looking to play triple-A blockbusters or genre-shifting indie titles, you'll find something of interest on page 48.

Our other big feature looks at where we are with solid-state drives. Prices have been tumbling to incredible levels recently, performance is better than ever, and as capacities increase, you can now legitimately build a PC that eschews spinning drives entirely. But there are still a few things you need to watch out for if you want to access the full potential of these drives, particularly when it comes to connecting several PCIe drives. Turn to page 36 to get the lowdown on what's what in the brave new world of SSDs, and also discover what our recommendations are for adding storage to your own machines.

Alan Dexter is Maximum PC's *executive editor and a punisher of hardware. He's been a tech journalist for over 20 years, and has no problem upsetting the PC industry as a whole.*

↘ submit your questions to: comments@maximumpc.com

THE NEWS

AMD Ups the Pressure on Intel

Zen's development cycle is still in overdrive

AMD CERTAINLY has ambition, and a hectic release schedule. It was only last March that the first Zen chips appeared, and since then we've had a myriad variants, and a second iteration of the design, the Zen+. The EPYC range of server processors came along in June, and repeated the trick Ryzen had managed on the desktop: It offered a price to performance balance that left Intel more than a little bruised. EPYC sold well, taking a decent bite out of Intel's Xeon market share, but there's much more to be had. We are about to see EPYC 2, codename "Rome," a 7nm server chip that is ready for sampling now, and due for commercial release this year.

EPYC 2 will be the first AMD chip to use the new Zen 2 architecture. This is a modular design, with 7nm process nodes, called "chiplets," built around a 14nm I/O unit, rather than Zen's more self-contained blocks. The architectural improvements include doubling the floating point width to 256-bit, and doubling the L3 cache, along with a slew of tweaks to branch prediction, bandwidths, pre-fetching, instruction cache, and more. Clock for clock, Zen 2 will be 10–15 percent faster than Zen+, minimum. The shrink also offers lower power consumption. Starting with a server chip shows where AMD's plans lie, and where a good part of the profits are.

Moving to a 7nm process means more room for cores. The EPYC 2 will double the core count, now 64 rather than 32, instantly giving twice the power per socket. It also gets PCIe 4.0 support, and will be socket compatible with the previous EPYC, and indeed the next one, "Milan." Performance? At AMD's Next Horizon event, an EPYC 2 was pitted against a dual Xeon Platinum 8180M—no slouch with a total of 56 cores. The comparison employed the ray-tracing benchmark C-Ray, and the prototype EPYC 2 box coasted to victory. A carefully selected demonstration, of course; the benchmark likes floating point and lots of cores, which helps.

The desktop version of Zen 2, the 3000 series, is all set for this year. Leaks have it that the first chips are eight-core, run at 4GHz, and use the same AM4 socket as the 2000 series. There will be 12 and 16-core versions to follow. There's also talk of a 32-core HEDT version. AMD's multicore performance is already outstanding—Zen 2 will push it even further.

The next target is single-thread performance—AMD wants to top the desktop gaming benchmarks, a prize that has eluded it so far, but Zen 2 could be the architecture to do it. Clock speeds are unlikely to jump by much, but the efficiency improvements could tip it over, unless Intel does something....

With AMD's blistering development cycle, Zen 3, an evolution of Zen 2, is promised in 2020. Zen 4, more of a redesign, is already being planned. Radeon is also going 7nm. First to emerge will be Vega-powered Instinct M160 and M150 cards. These PCIe 4.0 cards are aimed at machine learning in data centers. The M150 is due any day, with the M160 following before spring. 7nm graphics cards will follow, but details are scant. More precise details on the upcoming Navi cards are thin on the ground, too, apart from rumors that Radeon is to focus on consoles and servers first. We've yet to know much about any response to Nvidia's hardware ray-tracing Turing architecture either. A response is inevitable, but development teams are pretty busy right now, so it'll be later rather than sooner.

AMD appears to have skipped the 10nm process; the original Zen was a 14nm piece, the Zen+ optimized and shrank this to 12nm. The new Zen 2 will jump straight to 7nm for its processing units. This must be galling for Intel, which is struggling to get its own 10nm process off the ground, and is reduced to getting all the value it can out of 14nm chips.

It's only been 22 months since AMD got back in the game with the Zen-powered Ryzen, and the company has been firing on all cylinders ever since. EPYC 2 and its Zen 2 stablemates look set to increase the pressure on Intel across the marketplace, from servers down. **–CL**

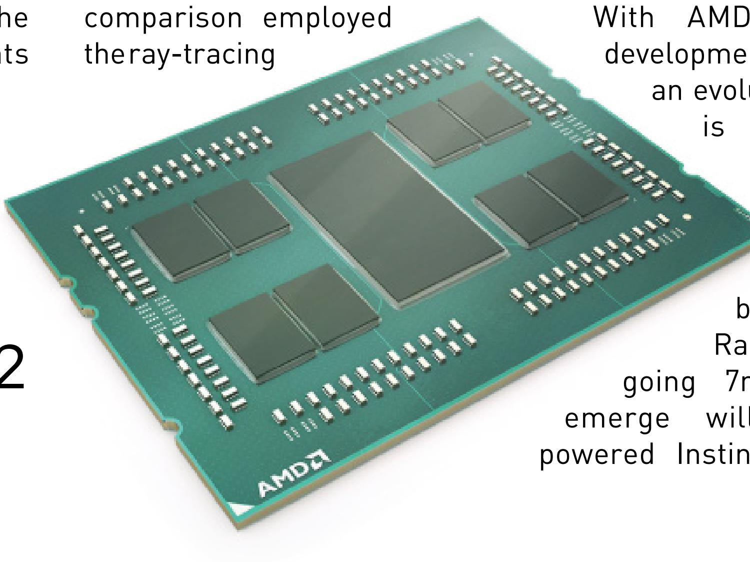

> “
> EPYC 2: the first chip to use the Zen 2 architecture

MAKE YOUR OWN PROJECTS WITH THE RASPBERRY PI

Learn the electronics, computing, and coding skills you need to make your own projects with the Raspberry Pi, and let your imagination run wild

NEW
✓ 100% unofficial ✓ Vital add-ons ✓ Python tips
Raspberry Pi
Tips, Tricks & Hacks
Code a Raspberry Pi synthesiser
Stream internet TV to your Pi
Raspberry Pi Tips, Tricks & Hacks Eighth Edition
100% UNOFFICIAL
ISBN 978-1-78389-645-5
EIGHTH EDITION
FUTURE

ON SALE NOW

Ordering is easy. Go online at:

www.myfavouritemagazines.com

Or get it from selected supermarkets, newsagents or bookstores

FACEBOOK TO FACE THE MUSIC

Internal document leak could reveal all

FACEBOOK HAS LURCHED from one public relations mess to another for a couple of years now. The latest and biggest problem stems from the poor way it managed news output during the 2016 elections. The story has now taken a bizarre twist. It is being sued by an app developer called Six4Three, whose owner, Ted Kramer, has called Facebook "the biggest violator of data misuse in the history of the software industry." Six4Three had an app that searched for pictures of people in bikinis on Facebook (not exactly a noble quest). Facebook shut it out and killed its business. As part of the case, Six4Three obtained private internal documents that revealed how much Facebook, and Zuckerberg in person, knew about the problems with privacy, and the abuse of it by Cambridge Analytica. These documents were placed under seal by the Californian court.

Now the bizarre bit: The documents were handed over to authorities in the UK while Mr. Kramer was visiting the country. Apparently, he panicked when approached on the matter on behalf of lawmakers investigating Facebook in the UK. Having tried to keep the documents private, it looks as though they'll be leaked abroad.

Meanwhile, Facebook is trying to sell its smart home display, the Facebook Portal; even heavily discounted, it's struggling. The idea of Facebook having a camera and microphone in your house could not have come at a worse time. It was already making the lists of the worst tech devices of the year in the month it was launched.

Zuckerberg has been making promises, and apologies, for years, but things haven't fundamentally changed. Trust is a valuable asset, and Facebook is running low. It is big, and very rich, so it won't implode, but we understand why it has hired an ex-diplomat as its head of its Global Affairs and Communication: It is going to need one to talk its way out of this. **–CL**

INTEL'S NEW TOP DOG

THE LATEST CORE i9 NOW COMES PACKING 18 CORES

INTEL HAS A NEW FLAGSHIP CPU: the Core i9-9980XE, a 14nm process Skylake-X that boasts 18 cores. There's an obvious rival in view, too: AMD's 32-core Threadripper. The Core i9-9980XE has a base clock of 3.0GHz, with a Turbo of 4.5GHz, and a useful all-core Turbo of 3.8GHz. Adding two cores over the i9-9960X has cost 100MHz on the base clock, but it receives a bump in L3 cache, now 24.75MB.

Performance is mixed. For straight single-core jobs, it loses to its sister chips, who can pump the clock better, especially the 5GHz available to the i9-9900K. In multicore jobs, it struggles to keep pace with Threadripper's finest. It's very quick everywhere, nicely balanced overall, but it's not often the absolute fastest anywhere.

Being Intel's king of the hill, you know there's a big number coming. The Core i9-9980XE launches at $1,979, $180 more than AMD's Threadripper, the 2990WX. For straight gaming, there are faster, cheaper Core i9s. It's lovely, but you pay for it. **–CL**

MICROSOFT TESTS ADS

All a mistake apparently

RECENT TESTERS of pilot versions of Windows 10 were surprised to find something extra in their inbox: ads. It didn't last long, the "feature" quickly disappeared. Microsoft has claimed it was an A/B test that was never meant to be seen by Windows Insiders.

The company's head of communication, Frank X. Shaw, said it was "an experimental feature that was never intended to be tested broadly, and is being turned off." If so, why was such a detailed FAQ prepared for users? An FAQ that clearly states that the company is running a pilot of the system in four countries to "get user feedback on ads in Outlook."

The ads appear in non-work accounts, including Outlook, Gmail, and Yahoo Mail. Their targeted nature can be turned off, via privacy settings, but this will still leave you viewing generic ads. If the idea has reached the point of A/B testing with working code, it's presumably at a fairly advanced stage. **–CL**

Tech Triumphs and Tragedies

A monthly snapshot of what's good and bad in tech

TRIUMPHS

LIFE IN MAGNETIC MEDIA

Seagate plans to push the HDD to 100TB by 2025, using new heat-assisted magnetic recording technology.

DDR5 ARRIVES

Hynix has produced the first JEDEC-compliant DDR5 memory modules, capable of 5,200Mb/s, 60 percent faster than DDR4.

MORE FLEXIBLE FUN

Samsung finally showed off its foldable phone at its Developer Conference. It launches in 2019.

TRAGEDIES

RIP STEAM LINK

After the current stocks are sold, that's it for Steam's Link; there wasn't any profit in the sales it generated.

THE COST OF MINING

A new report claims that mining cryptocurrencies uses more power than mining metal.

REPAIRS THAT BRICK

Apple has confirmed that its T2 security chip may kill your device unless you use authorized dealers for repairs.

Free VR Still Wouldn't Sell

"NO EXISTING OR IMMINENT VR HARDWARE is good enough to go truly mainstream, even at a price of $0.00." A stinging criticism of VR delivered by a cynical critic? No, this is from the blog of Palmer Luckey, the founder of Oculus VR.

"You could give a Rift and PC to every single person in the developed world for free, and the vast majority would cease to use it in a matter of weeks or months," he goes on to say. "Hardcore gamers and technology enthusiasts are entranced by the VR of today, as am I, but stickiness drops off steeply outside of that core demographic. Free is still not cheap enough for most people, because cost is not what holds them back...."

Barriers to adoption have previously been cited as high prices and inadequate technology, but both have been largely addressed—it's cheaper and better than ever. There are other more fundamental problems that can't be fixed so easily, though. The headsets can be cumbersome, and require your undivided attention. Configuration can be difficult, and no amount of marketing can build a tolerance to motion sickness.

No matter how impressive it becomes, some technology just won't be adopted by the masses. Witness the death of 3D TV, and the rise and fall of Microsoft's Kinect. Oculus's owner, Facebook, has recently moved the focus of the company to more mainstream models, which may be a bad move. Nobody is selling VR in big numbers.

Are Luckey's pronouncements a refreshing piece of realism? Possibly. VR has been "the next big thing" for so long, it's laughable. High-end VR is awesome, as Luckey says: "I hope and pray I am wrong, but most people are not like you and me." **–CL**

The founder of Oculus has doubts that VR will ever go far beyond hardcore gamers and enthusiasts.

The Slowdown of Linux

Meltdown and Spectre have cast a long shadow, not least on Linux. It has been a long, difficult process to secure the OS, and November's kernel 4.20 release returned to the issue, but this time the performance drag on Intel systems hit hard. Sorting out the Spectre variant 2 chip flaw means putting heavy restrictions on Hyper-Threading. Previous fixes knocked off a few performance percentage points; this new fix provides the security, but tests show a performance hit of 30–50 percent on many tasks. The effect has been particularly hard on Linux servers. The cure is close to being worse than the disease. Linus Torvald has even suggested turning off Hyper-Threading. **–CL**

Nvidia's Failing Flagship

Last month, we saw an alarming number of reports of Nvidia's RTX 2080 Ti Founders Edition cards failing. Nvidia has admitted that the "limited test escapes from early boards caused the issues." A test escape denotes a defect that has evaded all testing, a generous term to use—the early cards detonated soon after arrival, so clearly weren't thoroughly soak tested. There are no firm figures, but anecdotal evidence points to a failure rate of about half of the first batch. The standard RTX 2080 Ti has fared much better at about 2 percent. Still no word on what is failing, and why. Suspicion still centers on the memory, although a problem with the GPU itself hasn't been ruled out. **–CL**

FOOLING FINGERPRINTS

Print recognition proves to be fallible

A RESEARCH TEAM from Michigan and New York Universities has devised a method to break into fingerprint recognition systems. The remarkable, and alarming, thing is that the feat has been achieved without access to the original fingerprint or any hacking of code or hardware. It just uses a fake fingerprint that grants access.

These "DeepMasterPrints" are created using machine learning. An AI system was trained on thousands of prints until it could create realistic human fingerprints from scratch. These combine the most common traits of a fingerprint, producing a generic archetype, like a skeleton key.

Biometric security is convenient, and we've been told that our fingerprints and eyes are unique. However, the systems only work on a subset of the possible data; you don't need a full and perfect fingerprint impression to trigger access. They also only require a partial match of certain data points, so a partial print is used to partially match a stored one; close enough is good enough. Thus fakes can pass inspection, as enough features match a stored sample somewhere.

The DeepMasterPrint's success rate depends on the level of security of the device. It varies from under 1 percent on highly secure systems to a worrying 76 percent on the lowest settings. On commercial smartphones, it works on around one in five attempts. It's still at the proof-of-concept level—the researchers didn't actually hack any systems with conductive printouts—but it has raised a few eyebrows in the industry, which now has another problem to worry about. Tweaking the algorithms and bumping up the resolution and size of the readers will help, but it's worth remembering that a decent password remains the best way to secure your device. **–CL**

GET THE MOST FROM WINDOWS 10 WITH OUR COMPLETE GUIDE

All that you need to make the most of your Windows PC or laptop, including the best tips, tricks and tutorials!

NEW
Windows 10
Tips, Tricks & Apps
Unlock the full potential of Windows 10
Over 400 hints & tips inside
Learn how to customise your experience
Get more from the fantastic apps available
Organise your imagery in one simple place
✓ Personalise Windows
✓ Download apps
✓ 100% unofficial
Windows 10 Tips, Tricks & Apps Seventh Edition
ISBN 978-1-78389-612-7
SEVENTH EDITION
ON SALE NOW

Ordering is easy. Go online at:

www.myfavouritemagazines.com

Or get it from selected supermarkets, newsagents, or bookstores

Jarred Walton

TECH TALK

Battlefield V Ray Tracing

NVIDIA OFFICIALLY STARTED selling its new RTX 2080 Ti on September 20, 2018. Prior to the launch, we were given a short hands-on demonstration of ray tracing in a couple of games: *Battlefield V* and *Shadow of the Tomb Raider.* But both games ended up releasing without any form of ray tracing, with a promise that a future patch would add the feature. We still don't have a *Tomb Raider* patch, but *Battlefield V* made good on its promise, with a DXR support arriving just a week after the early access began.

That's the good news. The bad news: It's really demanding. We shouldn't be surprised, and it's not as bad as some suspected after the initial demos in August. There were rumors flying about that the game only ran at 30fps at 1080p with ray tracing, and that was on an RTX 2080 Ti. Reality is much better, with 60fps or more, even at 1440p on an RTX 2080—at least in some areas. More complex scenes can drop quite a bit from that mark.

As an example, an RTX 2080 Ti runs the snowy Nordlys single-player campaign at 92/70/41fps at 1080p/1440p/4K using the Ultra preset. Turn DXR off, and performance improves to 144/128/78fps. The RTX 2080 goes from 139/105/61 without DXR to 78/58/29fps with ray tracing, and the RTX 2070 drops from 121/90/52 to 65/47/25fps. Those are steep declines. Even worse, they appear to be more of a best-case result.

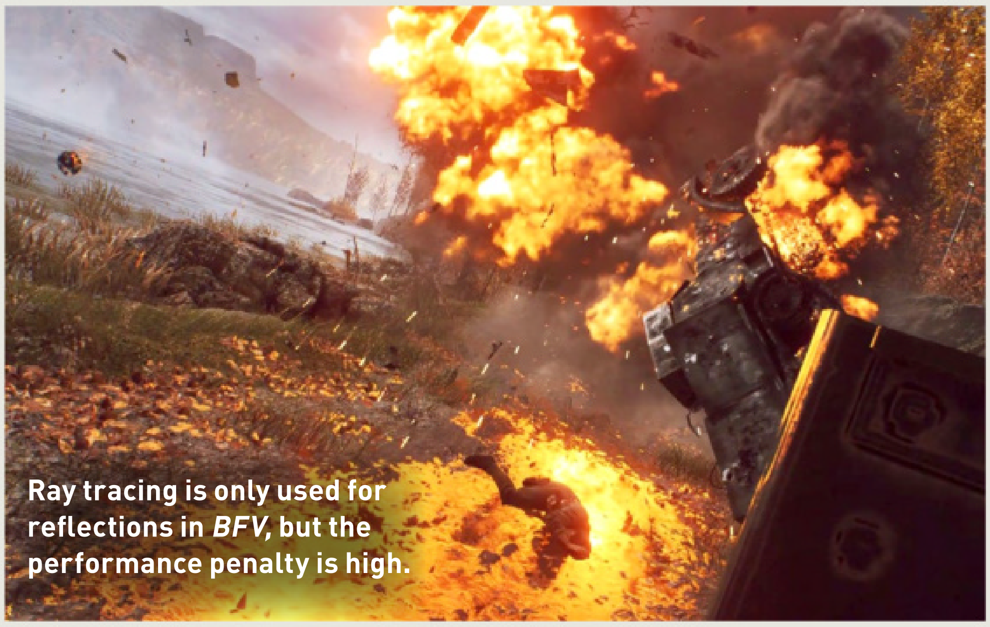

Ray tracing is only used for reflections in *BFV*, but the performance penalty is high.

On another level, Tirailleur, running without DXR gives slightly better performance than the Nordlys level I tested, but with DXR on, frame rates drop by 60 percent. That equates to 75/57/33fps on the 2080 Ti, 63/47/24 on the RTX 2080, and 50/38/20fps on the 2070. If that's not bad enough, performance in some of the multiplayer testing appears even worse, though the frequent complete stalls make me think there's still plenty of debugging left to do.

Frankly, *Battlefield V* is a poor initial showcase of ray-tracing tech. For one, DXR is only used for reflections—global illumination, ambient occlusion, and other effects still rely on the rasterization hardware. And while DXR ray-traced reflects look better than screen space reflections—which can only simulate reflections with data that's visible on the screen—it's not enough of a graphical upgrade to warrant the 40–60 percent drop in performance. It's also a fast-paced multiplayer game, so players are less likely to want to sacrifice frame rates in pursuit of higher image quality.

There's hope that things will get better, as DICE has stated that the current implementation of DXR is more of a first go, and it wastes a lot of computational effort casting rays that end up not being useful. You can also turn down the DXR level to reduce the performance hit—but, of course, that also reduces the visual impact.

This is still very early days for ray tracing, and I remain optimistic that other games will put the RTX hardware to better use. *Metro Exodus* and *Shadow of the Tomb Raider* are both primarily single-player, and they use DXR for more than just reflections. And to be fair, getting even one major game that uses completely new graphics technology within two months of availability is impressive.

My real concern is that this first round of RTX hardware is merely a stepping stone. It's a foot in the door for ray-tracing hardware, and it paves the way for faster and better implementations. The future RTX 30 and 40 series (or whatever they're called) are when ray tracing will hopefully go mainstream, and by then, today's $1,200 2080 Ti may perform more like a $200 piece of hardware.

Jarred Walton has been a PC and gaming enthusiast for over 30 years.

THE LIST

TOP 8 FREE WEB APPS JUST AS GOOD AS DESKTOP APPS

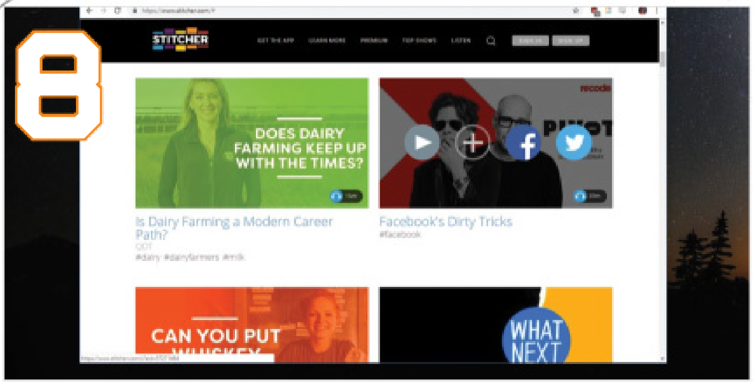
8
STITCHER
DOES DAIRY FARMING KEEP UP WITH THE TIMES?
Is Dairy Farming a Modern Career Path?
Facebook's Dirty Tricks
CAN YOU PUT
WHAT NEXT

STITCHER Why download podcasts? Listen online, and sync your listening with your phone. Magic!

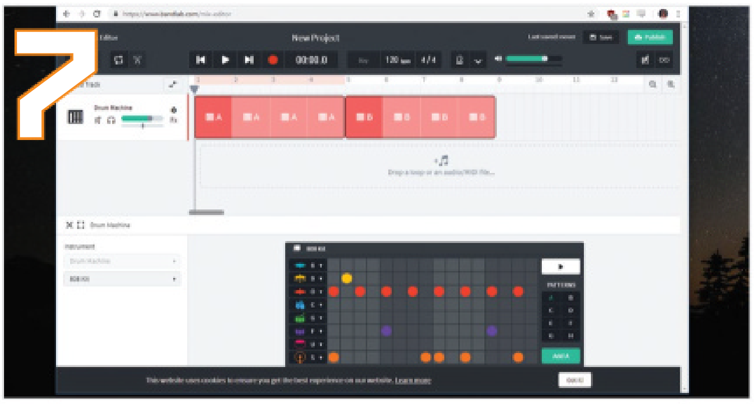
7
New Project

BANDLAB A free DAW that's actually functional, with the addition of integrated collaboration tools.

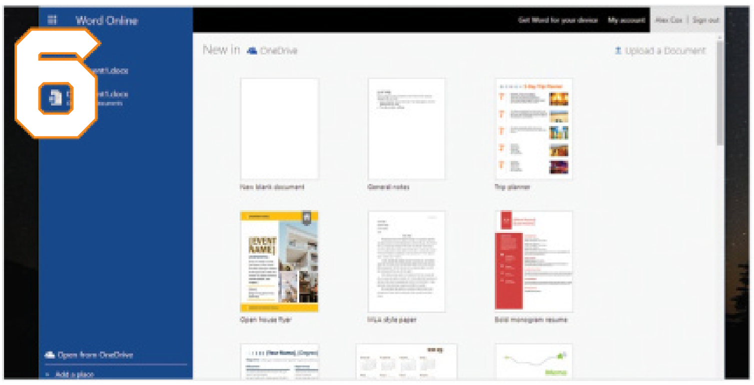
6
Word Online
New blank document
General notes
Trip planner

OFFICE ONLINE Better than full-fat Microsoft Office? Debatable. Significantly more free? Absolutely.

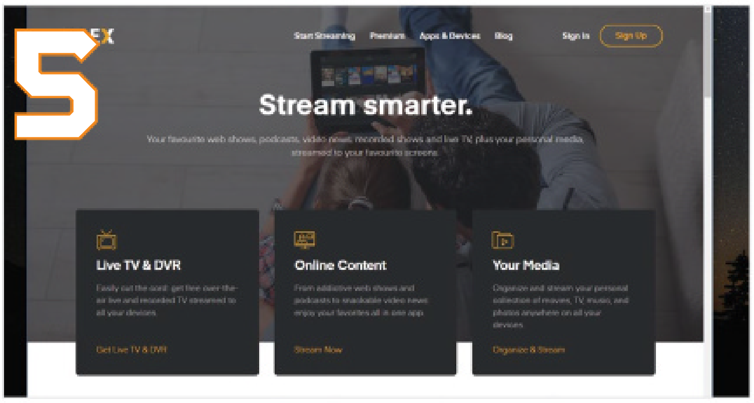
5
Stream smarter.
Live TV & DVR
Online Content
Your Media

PLEX Media Player? Pah. VLC? Forget it. Plex plays any content in any format, through a web interface.

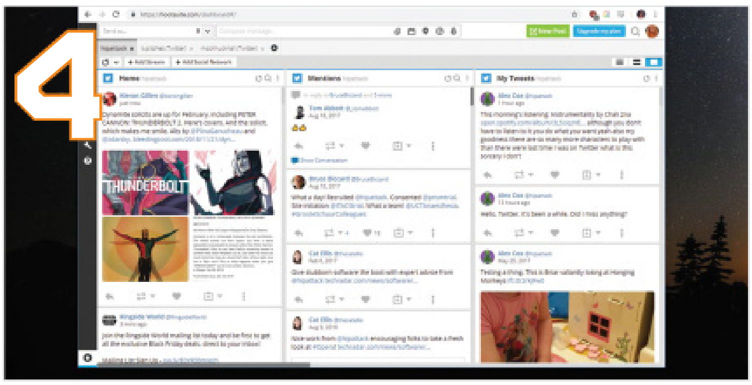
4
THUNDERBOLT

HOOTSUITE Support for Facebook and LinkedIn streams (plus Twitter) makes it better than Tweetdeck.

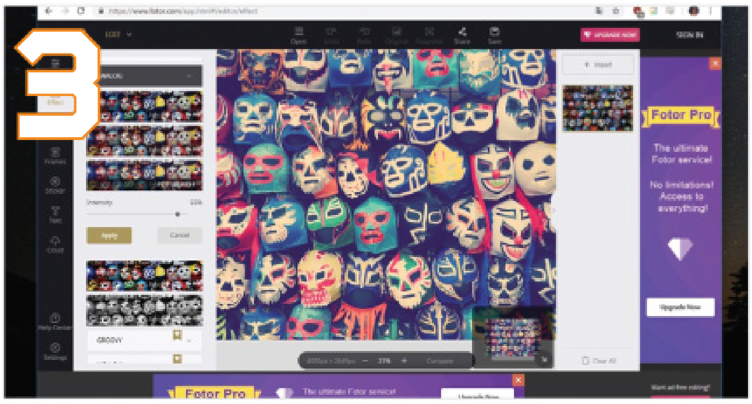

FOTOR Adobe Lightroom, get to stepping. This vast range of sexy filters can gussy up even the worst photo.

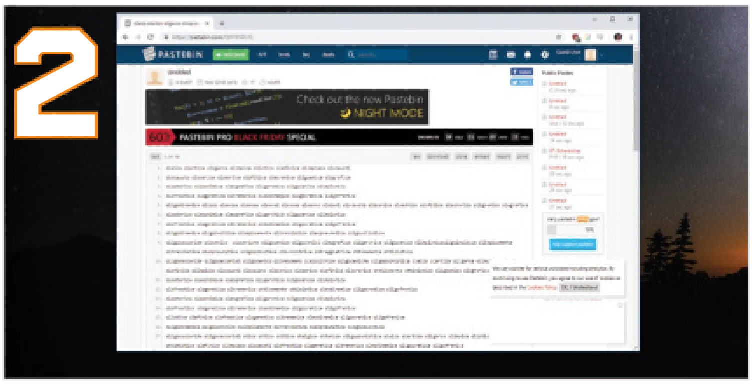
2
PASTEBIN
Check out the new Pastebin NIGHT MODE

PASTEBIN Why use Notepad when you could share your notes, or even set them to self destruct?

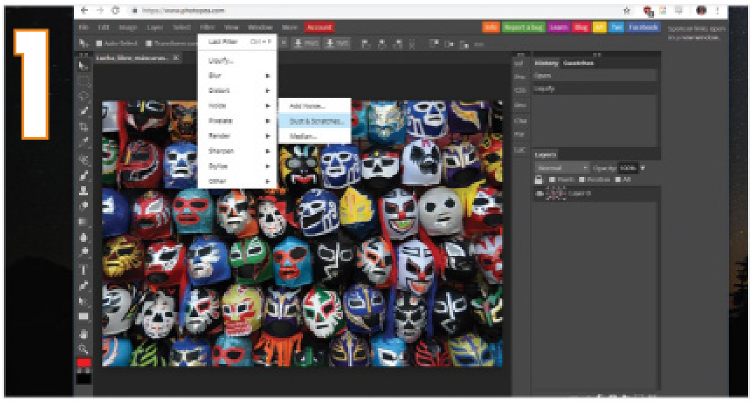

PHOTOPEA A Photoshop-alike that can open and edit more file types; there's even a good Liquify plugin.

Alex Campbell

OPEN SOURCE

You're Now in the Decade of Ubuntu 18.04 LTS

IN NOVEMBER, Canonical announced that Ubuntu 18.04 LTS will receive support for 10 years, doubling the lifespan of Ubuntu's historical long-term support releases. This seems crazy, but if you look at the Ubuntu ecosystem, specifically the server world, it's not so strange after all.

Consider first the timeline being discussed here. In the technology and software world, 10 years is an entire lifetime. Facebook got rolling in 2005, only 13 years ago. Instagram is only eight years old. Canonical itself has only been around since 2004.

A 10-year support horizon seems ambitious. But in a world where Canonical and Red Hat make their money catering to cloud customers, actions may seem puzzling to the desktop user. After all, Ubuntu has iterated through Linux kernels and provided new features via new version releases. Does that mean Ubuntu 18.04 will become stale?

To users of Ubuntu Server, a long-term support horizon of 10 years means more predictability. Applications will have a known stable platform to target. And applications for the server increasingly mean those that run in a Ubuntu Server virtual machine, or in a Docker container on top of Ubuntu, eliminating the reliance on the server to fetch packages from a repository.

Think about containerized applications, whether that be an application deployed as a Docker container, Snap, Flatpak, or AppImage. All those systems include the application's dependencies in the container. If an application needs a specific version of Python 2.7, for example, there's no need to ensure the system's Python version is correct. The user can install the containerized version, which comes with its own copy of that Python version. If that's the paradigm, what's the value in retiring an older version of the base OS? The libraries running on the OS become increasingly irrelevant at the application level, relegating OS updates to security patches for the base system and the occasional kernel upgrade.

For the desktop, the story is becoming much the same. Canonical has been pushing Snaps on Ubuntu users for a while. The GNOME Software app store includes Snappy packages as well as packages available through Apt repositories. (Fedora does the same, but with Flatpak instead of Snappy.) AppImage is also gaining popularity among developers as a generic format. In fact, one of the big desktop updates in Ubuntu 18.10 was better GUI integration of Snap package administration.

Does this mean an end to Ubuntu versions? There's no indication (yet) that this is the case. Ubuntu has released LTS versions in the April of every even year since 2012. Ubuntu 14.04 is still receiving support until April 2019, for example. LTS releases are hardly the only releases that users have available. Ubuntu also releases versions in the April of odd years, as well as annual versions around October. While much shorter-lived, these act as testing grounds for where the OS is moving, and can be a preview of things to come. (Ubuntu dumped the Unity desktop for GNOME in 17.10 before the 18.04 LTS release.) I still expect to see Ubuntu 19.04 come April.

As with the IBM acquisition of Red Hat, these moves by Ubuntu will probably not have earth-shaking effects on desktop users. Canonical, like IBM and Red Hat, is catering to the server and cloud customers who pay the bills, even if they claim to still love the desktop.

In a post on the Ubuntu blog, Canonical's James Nunns said, "Whilst Canonical has lots of projects, we are not shifting away from Ubuntu Desktop—we love it and we love that users love it. Canonical and Ubuntu are intertwined, the success of one supports the other and vice versa."

People do love Ubuntu and derivatives like Mint as a desktop OS. But it's obvious that Canonical's development efforts in projects like OpenStack will outpace those of the desktop. I'm not saying this is a bad thing, but I miss the days when every new Ubuntu version meant an awesome new desktop feature or improvement.

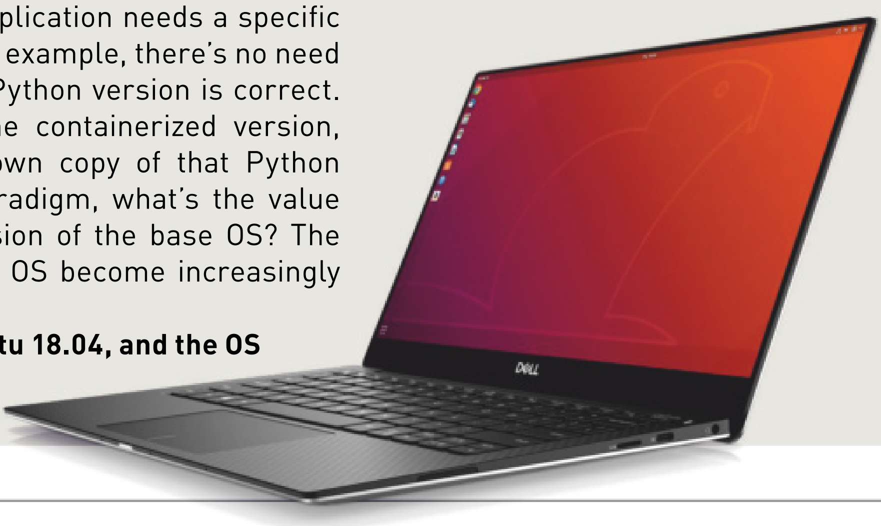

Buy an XPS 13 with Ubuntu 18.04, and the OS will be OK for a decade.

Alex Campbell is a Linux geek who enjoys learning about computer security.

HEAD TO

BY **ALEX COX**

Smartphones vs. Tablets vs. Chromebooks

You need the Internet, you need to get things done, and you need it in a convenient form, but what form does that wonder device take? Is it a tiny sliver of tech primed to explode in your pocket? A slightly bigger sliver of tech, ready to go bang in your bag? Or something laptop-shaped, but only partially laptop-ish? Bashing together three very different things is always difficult, but each category fills that connectivity niche nicely—and needn't break the bank.

ROUND 1

Productivity

This won't come as a surprise, but it's very hard to work efficiently on a smartphone. Even the largest phablets have a screen that's just too small to contain any significant portion of a document and an on-screen keyboard, and while the apps are there—analogs of the Microsoft Office suite, Google's collection, and everything in between—actually using them is another matter. The same can be said of tablets; if you really want to get something done on a keyboard-free device, you're going to have to juggle with Bluetooth peripherals and squint really hard. That said, there's slightly more screen real estate to work with on bigger devices, and certain examples (the iPad Pro, in particular) are uniquely suited to artistic pursuits in a way that no other device can match.

While the tools on a Chromebook are more limited than those of a full PC, and the power is far less, the laptop-shaped form factor lends itself much more sharply to getting things done. Just about every task can be completed with web-based apps, and (rickety) compatibility with Android apps means you'll likely be able to fill in any blanks if you run into a wall.

Winner:
Chromebook

ROUND 2

Media Consumption

It's hard to split our categories, because they each have their own advantages. A Chromebook can be placed on a surface, angled appropriately, and chuck out as many movies and online videos as you like while you lounge. But we'd have to list it bottom of the pile because, unless you invest a lot of cash in a Chromebook, you're going to be dealing with a barrel-bottom display—and throwing a thousand bucks at a Chromebook basically defeats the object.

Tablets give you that larger screen, and usually sport solid viewing angles—there's nothing better (apart, perhaps, from a massive TV) when it comes to curling up on the couch to watch something special. Load up a tablet with media before you travel, and it helps you pass the time in style. But is a tablet as strong as an LTE-equipped phone that can slide into your pocket? Not by a long shot. Podcasts, audiobooks, streaming radio—they're all best suited to a tiny device with headphones, and the negative of a phone's smaller screen doesn't beat the positive of its portability. Mobile media consumption on either of the other formats just isn't as convenient.

Winner:
Smartphone

ROUND 3

Accessibility

There are two meanings of the word "accessibility," and one is a synonym of "availability." That means there are two outcomes to this category, and two ways to judge our devices. Think about the phone, for example, that ever-present Internet-connected computer that's always there when you need to look at cat pictures or dodge robocalls. Top of the tree, perhaps, but do a little experiment: Pull your phone out of your pocket, open your messaging app, close your eyes, and write a message. Attempting the same trick with the tactile keyboard of a Chromebook will likely net you far better results.

If you need the assistance of accessibility features, you'll find them on just about every device worth its salt. These can vary a little depending on the exact hardware, but you'll almost certainly find high-contrast graphics modes, magnifiers, screen readers or select-to-speak functions, and dictation functionality on any model of any of these devices. The scales, really, are wobbling between the phone and the Chromebook, and the latter has more weight—that keyboard scores it extra points, and adds a sprinkling of additional disability aids along with it.

Winner:
Chromebook

HEAD

Razer's entry into the smartphone market blurs the lines between devices.

Can a good old-fashioned tablet still keep up?

Chromebooks are affordable, versatile, and pack a punch.

ROUND 4

Performance

Again, this is a topic that depends on the precise hardware; each category has its low and very high-end devices. If you're spending $60 on a dinky Amazon Fire tablet, you can expect some awful performance and sluggish response times. If you drop upward of $1,000 on an iPad Pro, you'll get a far more sprightly interface. Cheaper phones and cheaper Chromebooks have the same problem; they're underpowered, built to hit a mark, and look nice in pictures. Lower-resolution screens, lesser processors, and skimping on RAM makes a big difference.

So, let's look at the higher end. Benchmarks don't help, as we're breaking down three separate worlds, making this more a holistic assessment. There is no blowaway Chromebook. The Pixel Slate, for all its bravado, is not a blowaway tablet, and even flagship Android phones have their niggles. The latest-generation iOS devices are—however you feel about their creators or their philosophy—by far the most responsive and pleasing slabs of mobile hardware, and for all the tightness of the iPhone, the A12X Bionic inside the 2018 iPad Pro is the processor to beat.

**Winner:
Tablet**

ROUND 5

Price

Goodness, aren't phones expensive? Yes they are. And progressively more so. Not all of them—there's a constant bleed of fresh brands from China that seem to be pushing the boundaries without pushing their luck on your wallet—but, in general, phones are far more expensive than they should be because they're status objects first and functional second. We're almost not angry about it, either; that cash fuels a drive for innovation that is pushing technology forward.

The real contenders here are the other pair, and we need to make a judgment call. Tablets can be had very, very cheap, but Chromebooks come in at astonishing prices, given that they're essentially proper computers, at least if you ignore the hamstrung OS. Those devices that do hit that bottom line aren't always the best, but as we'd be far more likely to recommend an 11-inch Chromebook over a rickety 8-inch Android tablet, there's only one sensible choice. Do note, though, that shopping around at the right time can bag you full-bore Windows laptops for the same price as a low-end Chromebook—and that's a license to do more.

**Winner:
Chromebook**

And the Winner Is...

The Chromebook is a great idea. We applaud Google's moxie in pushing forward a device with lower hardware standards and an OS that all but forces you online. It's a great way to tie people to Google's services, yes, but recent developments have made Chrome OS more flexible and viable as a standalone device. It's won the most categories in this subjective test, but does it win overall? No.

Neither does the tablet bag the crown because, after all this time, there's only really one good brand of tablet to buy; and while we've avoided talking about PC tablets, we'd lump most of them in the same slightly wobbly category as Android tablets. Nobody's quite got it right after all this time, and PC tablets are generally just laptops in disguise anyway. The smartphone is our winner. You need a phone. You have a phone. That phone fits in your pocket, connects you to the sum of the world's knowledge, and keeps you entertained hour after hour, wherever you may be. As much as they try, the other two categories can't replicate that killer feature. ⏻

DOCTOR

THIS MONTH THE DOCTOR TACKLES...

> Adapting to M.2
> Storage Solutions
> DirectX Drama

Corsair's Neutron NX500 drops into a four-lane PCIe 3.0 slot to deliver 3GB/s sequential write performance.

NVMe on X79

Hi, Doc. Ever since I built my current rig back in January of 2014, I've watched SSDs bump up against the limitations of SATA 6Gb/s, and then be freed once again by NVMe. The problem is, most NVMe-based SSDs are designed to drop into an M.2 slot, which my Asus X79-Deluxe motherboard lacks. I see that there are a few PCIe interface options, but I also noticed a number of PCIe-to-M.2 cards available. These would open the world of M.2-based drives to me, if they work. My question is: Do you have any information or caveats on the adapter cards out there? Are there any manufacturers to steer clear of? **–C. Steinway**

THE DOCTOR RESPONDS: First things first: Make sure that you download the latest firmware for your X79-Deluxe motherboard before installing an NVMe SSD.

With that base covered, you shouldn't have any trouble upgrading your storage subsystem. Add-in cards are indeed rarer than NVMe SSDs with M.2 connectors, but they do exist. Take a look at Corsair's Neutron NX500 or Plextor's M9Pe, for instance. If neither of those interest you, a PCIe-to-M.2 adapter should be fine. Its only purpose is converting physical interfaces. Four lanes of PCIe 3.0 from an expansion slot become four lanes of PCIe 3.0 through an M.2 slot.

HTPC vs. Smart TV

Doc, I own several TVs and Blu-ray players, and I notice that their "smartness" varies from one model to the next. The more modern (and expensive) consumer electronics naturally have more so-called intelligence.

Recently, I was thinking about buying a new top-end DVD player for about $350, because I have trouble accessing my Asustor Media Server, where I keep family movies, pictures, some music, and other audio stuff. I also wanted to prep my media cabinet for the purchase of a new 4K TV.

But then I had the idea that it might be better to keep a PC next to the TV. After all, it would have more "smartness" than any DVD player or TV with an embedded operating system. If the PC had an HDMI output and a decent graphics card, it'd serve my purpose nicely.

Do you think I can build a small PC to serve as a media server for around $500? It would be a big bonus if this setup could capture some TV content, too. I've never used a DVR, and I generally don't worry about saving anything on TV, but I might be tempted to experiment with this.

Any thoughts? If you believe this is a worthwhile endeavor, then please suggest some hardware that would work well in the context of a home theater PC. I don't need to spend big bucks on a graphics card designed for high-end gaming; it just needs to be good enough to drive a 4K TV. **–John Gaydos**

THE DOCTOR RESPONDS: Yes, it is possible to build a fully functional HTPC for $500 or less, particularly if gaming isn't one of your priorities.

In the Doc's house, there's an Intel Next Unit of Computing (NUC) hooked up to his stereo receiver via HDMI. The compact system (and its big SSD) host an entire library of audio and video. It's also responsible for streaming Hulu, Netflix, and HBO subscriptions. It runs Jackbox during game nights with friends, Pandora for background music, and Lexia Learning when the kids get home from school.

If you want to add over-the-air digital TV to the mix, check

↘ submit your questions to: doctor@maximumpc.com out SiliconDust's HDHomeRun DVR subscription, which sells for $35 per year. The software records from any compatible HDHomeRun tuner to any storage path you assign, and there's no limit to the number of devices that can watch recorded content simultaneously. The HDHomeRun Connect Cordcutter bundle, which includes two OTA tuners and an indoor/outdoor antenna for $100, would be a good place to start if you don't require access to premium channels.

Storage Conundrum

Doctor, I need help sifting through all the various storage devices out there. I am an avid photographer and videographer, and I use my PC mostly for video and audio editing. Right now, I'm running an Intel Core i7-3770K, with 32GB of DDR3 RAM, and an Asus GeForce GTX 1060 6GB. I know the hardware is a little outdated for video work, but it's fast enough for now. I am slowly upgrading the internals when I can.

Beyond this PC's performance, my main concern is that something bad will happen and I'll lose all my saved data. What are some good options for safeguarding my video projects that won't break the bank, yet still offer plenty of capacity to grow? Do you have any opinions on what's better between a standalone external hard drive or a network-attached storage appliance? Finally, if I go with a NAS, what is the best RAID configuration for me to use?

–Todd Palmer

THE DOCTOR RESPONDS: If your projects are currently saved to a single disk inside your PC, they are sitting ducks. A hard drive failure, a Windows update gone bad, or even an errant couple of clicks could wipe them out entirely. It's good that you're thinking about backup now, rather than after disaster strikes (and recovery options become a lot more expensive).

As a rule of thumb, maintain three copies of your important data. Two can be local, but one should be saved somewhere offsite. Adding an external drive or NAS appliance that Windows backs up to automatically would give you a second copy. Then, upload those irreplaceable files to the cloud using an unlimited data plan from Backblaze or Carbonite. In the unfortunate event of a fire or burglary, you're still covered.

If the Doc has to choose between an external drive connected through USB or eSATA and a network-attached appliance, he picks the NAS for its autonomy. Many, many years ago, he configured his five-bay appliance in RAID 5 with a hot spare. That's the way it operates to this very day. However, the Doc also understands the main drawback of RAID 5: A single drive failure makes the array vulnerable until it's rebuilt. Nowadays, it takes an inordinate amount of time to perform parity calculations on many-terabyte disks, so RAID 5 is far less attractive.

RAID 6 is proving to be a popular alternative in five-bay appliances as a result. A second parity stripe greatly increases the complexity of those aforementioned calculations. However, if one disk fails, the array is still running as an equivalent RAID 5 and can technically survive a second failure. You do give up two drives' worth of capacity for the distributed parity data, so make sure the disks in your appliance are large enough to compensate.

Falling Frame Rates

The HDHomeRun Connect device and DVR service stream can record OTA digital TV.

Dear Doctor, I'm having a rather strange problem with an MMO game *(Lord of the Rings Online)* that I like to play. Starting two or three months ago, my frame rates would drop to nothing for about 15 seconds, then go back to normal. This occurred every seven minutes. Needless to say, it made a game that I really enjoy almost unplayable.

I contacted the game masters about my problem, and they have been as helpful as they can be. I ran the diagnostic tests they recommended, including MSInfo, DxDiag, WinMTR, and so on. Plus, I updated the drivers for my Realtek audio device and GeForce GTX 960.

After updating the Realtek drivers to the latest version, there was a period of about two weeks when the lag problem disappeared. Then it came back and hasn't gone away again. It's almost like something is hogging the CPU for 15 seconds at a time and then relenting.

I have the latest version of DirectX 12 installed, but I still get this message when I start up my computer: "There was a problem initializing DirectX 9. Please ensure your video drivers are up to date and you have the latest version of DirectX installed." I'm not sure how to install an older version of DirectX, or why I'd need to do that if I already have the latest version installed. Are older iterations not supported by the latest version? This is very confusing, especially as this may not have anything to do with my lag problems.

I also ran a program called Core Temp with *Lord of the Rings Online* running as well, and it shows that my temperatures are normal (38–45 C). Any insight you might be able to pass along would be much appreciated. I'm not tech-savvy and only know enough about computers to be dangerous, so explanations would be helpful. I do enjoy reading *Maximum PC*, though! For what it's worth, I'm using an Asus Z97-A/USB 3.1 motherboard and Haswell-based CPU. I also have 16GB of Kingston DDR3 memory.

–Rob Duncan

THE DOCTOR RESPONDS: It's interesting that updating your audio drivers rectified the performance issue for a couple of weeks, so try recreating the conditions that led to a solution the first time. Press the Windows key, open "Settings," and click "Apps." Scroll through your list of "Apps & features," and uninstall the Realtek and Nvidia drivers. Visit www.realtek.com and www.nvidia.com for the latest versions, and install them. In the case of Nvidia's driver, choose the "Custom (Advanced)" option, and select "Perform a clean installation" to reset the driver's settings and clean out any old profiles.

Just to be safe, the next time you shut down your PC, open the chassis and check to make sure it isn't loaded with dust. You mentioned normal CPU temperatures, but it's important to check the graphics card's cooler as well.

The Doc also looked into your DxDiag log and noticed crashes in CompatTelRunner.exe. This may not be related to your slowdowns, but because Windows Compatibility Telemetry sends usage and performance data to Microsoft, it might be worth disabling if your hard drive is getting hammered. Open the Task Scheduler, expand the "Microsoft" tree, then the "Windows" branch, and click "Application Experience." For each task listed, right-click, and select "Disable."

Finally, the DirectX 9 error message you listed appears to be related to FRAPS. If you have the benchmarking/recording utility installed, consider uninstalling it or preventing it from starting with Windows. Hopefully, one of those steps addresses the root cause of your performance degradation.

2019 TECH PREVIEW

Take your mind back a moment. It's November 1996, Bill Clinton has just beaten Bob Dole to the presidency, "Macarena" is topping the charts, and *Star Trek: First Contact* is about to make waves at the box office. In the world of tech, however, something more wonderful is about to happen. Andrew Grove, Intel's CEO of the day, is about to drop one of his truth bombs, and that truth? Well, it's that by 2011, thanks to a combination of Moore's Law and genius architectural advancements, Intel CPUs will be operating at the 10GHz mark and beyond.

Although that prediction may have fallen a bit flat over the years, it's not to say we haven't advanced since then, technologically at least. The world of the transistor stops for no nerd, and as another orbit of the sun comes to a close, it's time for us here at the Maximum PC Lab to condense all of our prophetic ability into one clean, crisp, potentially slightly wooly, article. So, folks, likely more accurate than your election-time pollsters, and definitely more truthful than those Hurricane Florence weather forecasts, hold on to your hats, because it's time for *Maximum PC's* Tech Preview 2019. –ZAK STOREY & JEREMY LAIRD

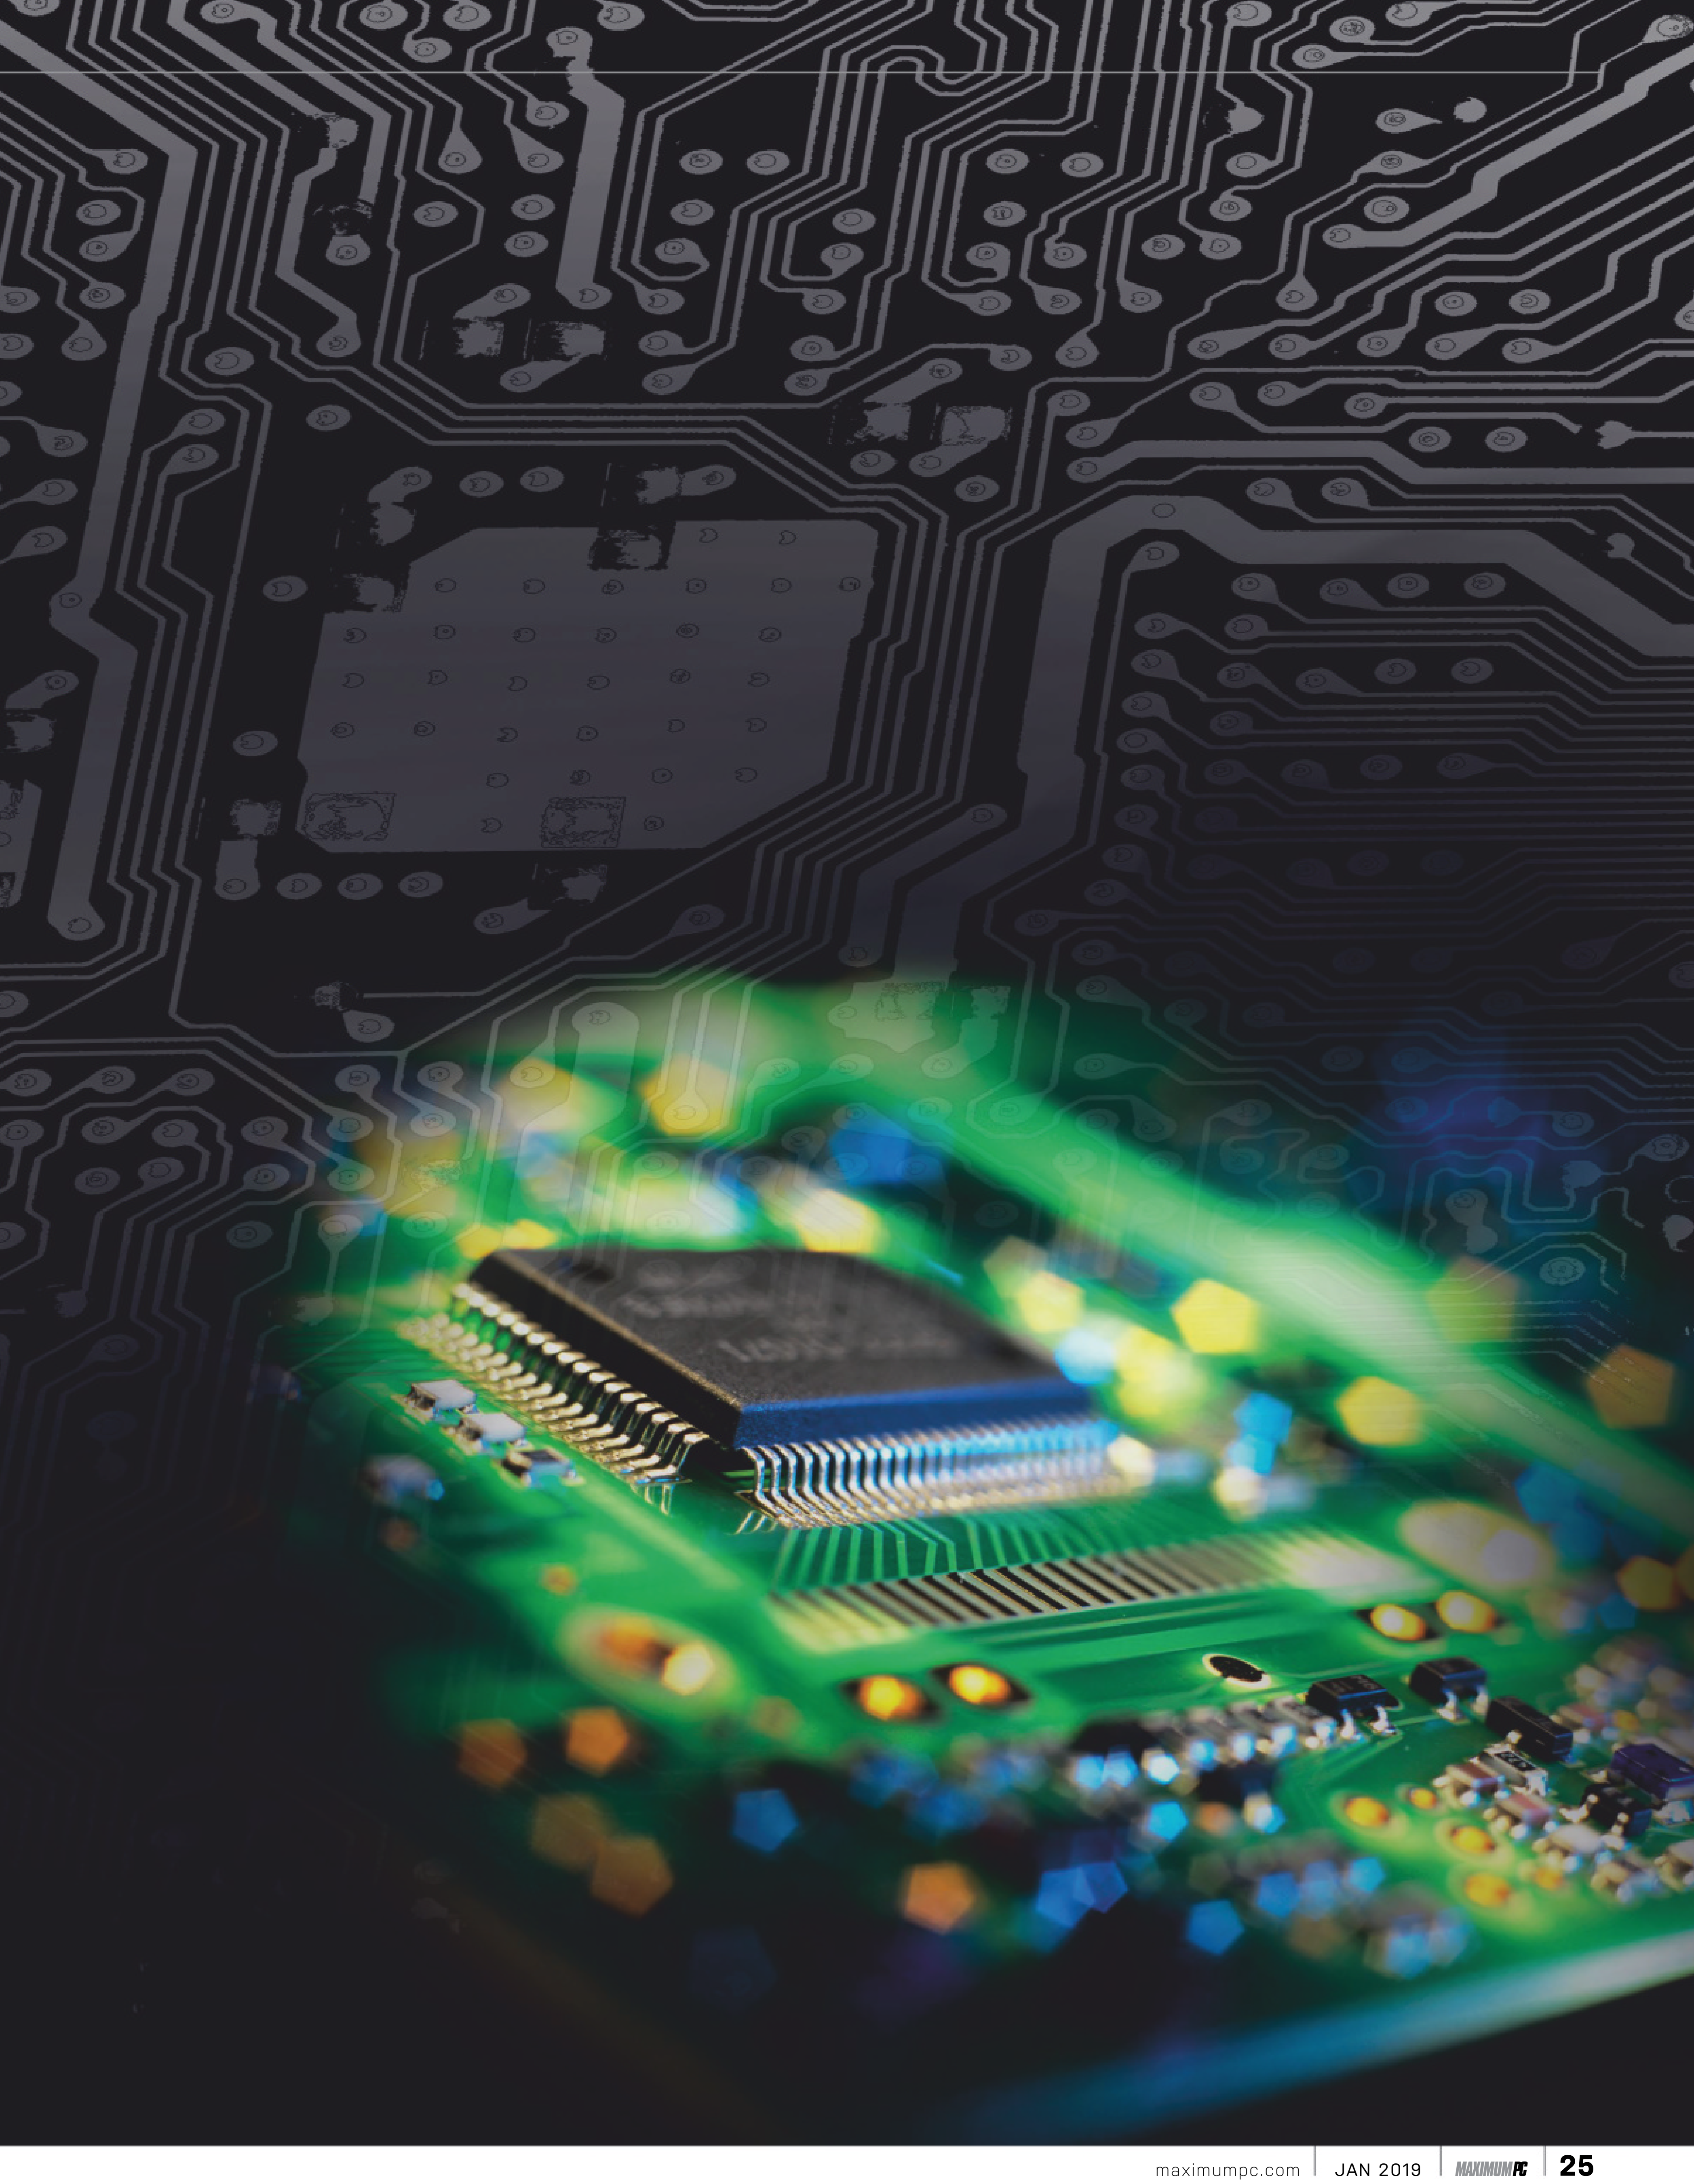

THE BATTLE FOR PROCESSING SUPERIORITY

What does the future of high-end tell us about the destiny of mainstream processing?

NOTHING'S BETTER for the tech industry than a little competition, and that's something Ryzen has very much brought to the table over the last two years. Consider this: Since the beginning of 2017, we've doubled the number of cores available to mainstream processing parts. From the very bottom rung, with Pentium and Athlon CPUs, all the way up to the Core i9s and Ryzen 7s, Intel and AMD have produced an epic lineup of chips, with twice as many threads as before. Look at high-end desktop processors, and that figure has almost quadrupled. Crazy. In early 2017, Intel's 10-core Core i7-6950X would set you back $1,723; today, you can buy AMD's TR 2990WX, with an incredible 32 cores, for just $6 more.

Are we likely to stop here? We doubt it. With AMD announcing its EPYC 2, Rome, server processors, due to launch in the second quarter of this year, the core war is still very much in full flow. Packing 64 cores and 128 threads on a single processor, with a far superior memory interface design than previously, supporting PCIe 4.0, and on TSMC's 7nm manufacturing process, AMD's next-generation server parts look as though they're going to be phenomenal computational monsters. There are even rumors of a 29.4 percent performance increase in single-core IPC, all down to architectural advancements and floating point calculation optimizations, rather than simply increased clock speed and voltage.

EPYC 2 is important for a number of reasons. Firstly, because it runs off the successor to the ridiculously good Zen architecture, the 7nm Zen 2, which will also be found in the third-generation Ryzen parts (also due this year). And secondly, as the first-gen Threadripper processors were essentially cut-down variants of the first EPYC server parts, it gives a good guess as to what Threadripper 3 might look like, too.

The big question is: What's going to transpose down? Apart from eight-channel memory and the PCIe 4.0 enhancements, which we're hoping make it, what will be more intriguing is how many cores are brought along with it. We imagine AMD will likely keep the overall design of the two SKUs the same, to help reduce manufacturing costs, but as we've already seen 32 cores on the mainstream Threadripper platform, it does leave questions as to how many cores will make the jump. At a guess, we may see a high-end 3995WX, likely the same price as Intel's unreleased 28-core HEDT part, potentially touting 48 cores and 96 threads, for the true amateur renderers and CAD-CAMmers.

We also know this is the first time AMD's deviating away from a single-size transistor design across the chip. Thanks to its Infinity Fabric, the company will be going to a modular transistor size architecture design, with things such as the ridiculously fat I/O controller staying on 14nm, while the main core complexes will be on 7nm.

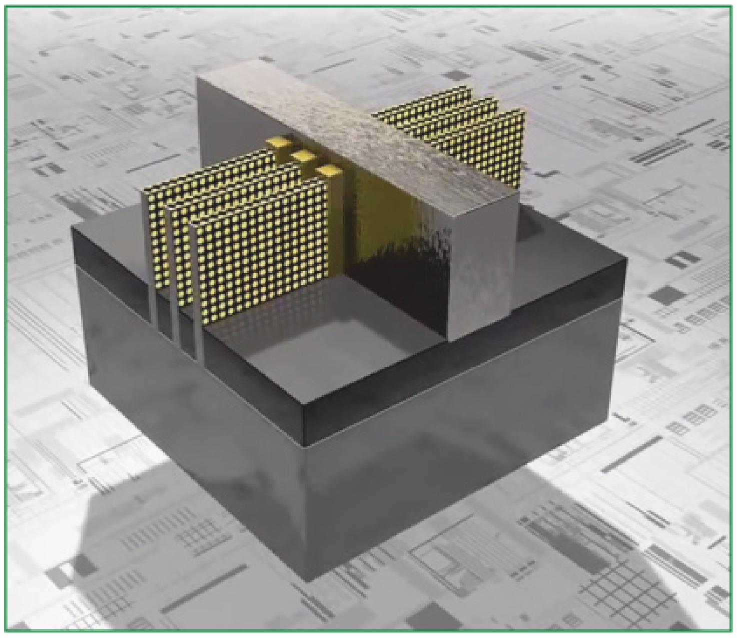

A three-dimensional representation of a tri-gate 14nm transistor.

Hopefully, most of that should filter down to Ryzen 3. We'd love to see a 16-core Ryzen 7 3700X, with support for PCIe 4.0, and perhaps more PCIe lanes in general for better PCIe SSD support, hooked up in a similar manner to its Threadripper kin—16 for GPUs, and an additional eight for PCIe storage would be a fantastic solution, and really future-proof you as PCIe SSDs break that 3.2GB/s barrier.

THE QUEST FOR 10NM

If the reported figures are correct, and AMD's samples of the EPYC server parts prove as mighty as they currently are, Intel may lose the top spot in multi and single-threaded applications. That would be the first

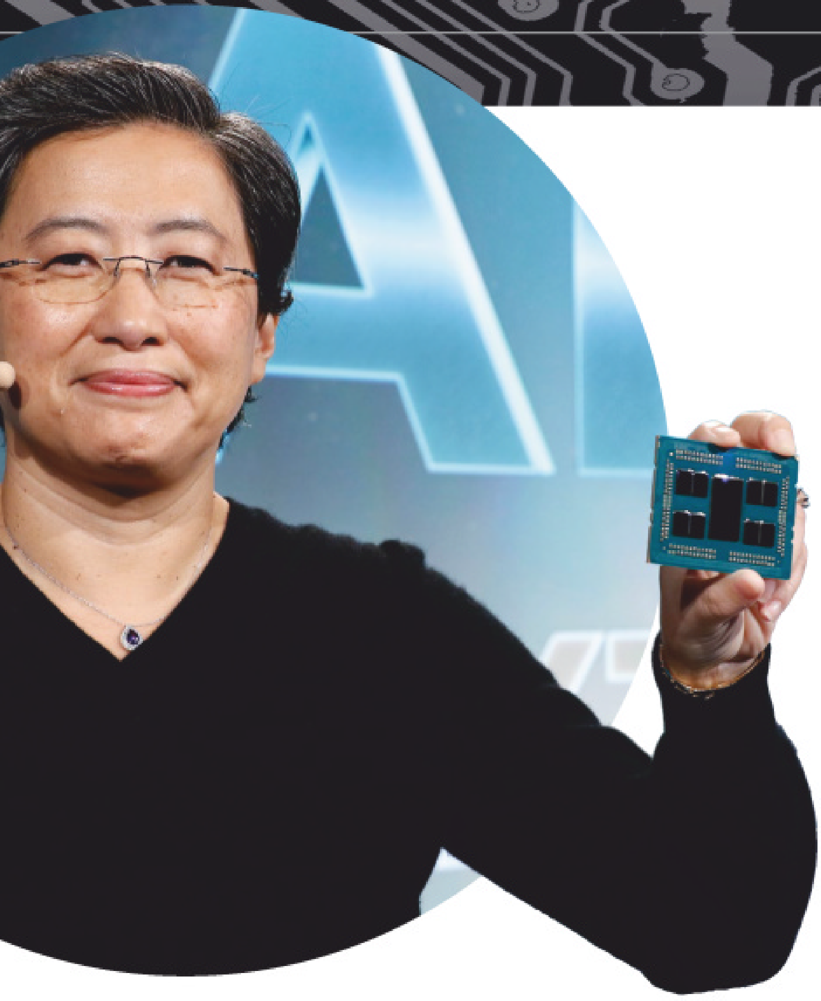

Welcome to Rome!

time since the early Athlon days that Intel has lost its grip on the top spot of processing performance.

It all rests on whether Intel can hit its 10nm volume production targets in its current timeline. Although originally outlined to make it to market by 2016, delays and complications have lead to multiple refreshes of Team Blue's 14nm architectures, with its current roadmap suggesting the company should be able to make bank on its volume production targets by the second half of 2019, leading to a late fall release at the earliest. That said, there is one 10nm processor currently in existence: the Core i3-8121U. It's an ultra low-power mobile part by design, but in a rather bizarre twist, the integrated GPU is switched off. According to sources, this is due to logic incompatibilities between the GPU and CPU that are currently too troublesome to rectify.

Going from 14nm to 10nm should allow Intel to increase its transistor density by a factor of 2.7. However, as in the past, this likely won't lead to a huge performance increase, so far as IPC is concerned, and with no evidence or rumors on benchmark performance, we can't say how powerful these parts will be. All of this leaves AMD with eight to nine months of open market to capitalize on, if it can successfully get its 7nm processors out of the door in time.

THE FUTURE OF PROCESSORS

With all of this coming from the mainstream manufacturers, you might be forgiven for thinking that's it for the world of the humble processor. Well, there's one outlier you might not expect to be making waves, but there are more than just ripples coming from our long-misjudged technical brothers across the battle-ravaged OS ocean. Apple is seemingly picking up the pace when it comes to processor design and overall performance. Take a look at the ARM-based Bionic A12 X processor found in its latest tablets, and even in the iPhone XS, and you'll see performance numbers equal to a mid-clocked Sky Lake Core i5-6400 or so.

That might seem like small fry when you look at the current state of desktop processors, but given the fact that the company has been working on its own bionic processors for a little over six years, coupled with Intel's continued delays in hitting its transistor targets, and the big fruit might finally be making its way to forging its own processors to replace the Intel ones found within its consumer devices. This is further backed up by a report from Bloomberg, back in 2017, stating that Apple plans to replace Intel's processors inside its Mac lineup by 2020.

Given that the iPhone XR scores a hefty 4,793 points in GeekBench's single-core test, in comparison to an AMD Ryzen Threadripper's 4,801, it's not hard to see where this could go if Apple really pushes the boundaries of that ARM processing tech.

Apple's got its foot in the door of single-core performance.

APL1W81 339S00560
A12
D9WCL
R9 07K 1820
L1H1Y81446 1833

THE GRAPHICS OF TOMORROW

AMD looks to the mid-range; Nvidia still holds the high ground

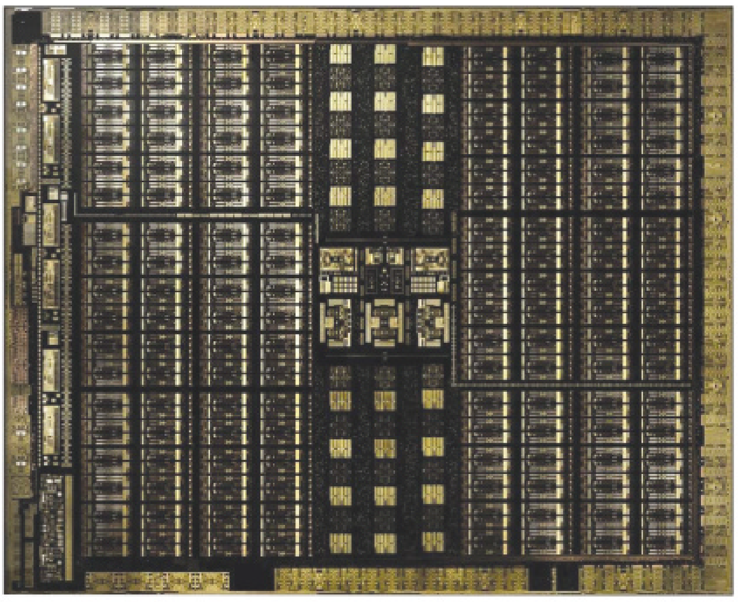

A full-fat Turing die.

2018 MAY HAVE SEEN the advent of ray-tracing for Nvidia, but Turing actually left a lot to be desired. The full-fat architecture provided a solid 30 percent performance increase, but as it cost over 70 percent more—thanks to the introduction of dedicated ray-tracing and DLSS hardware—that boost felt more than a little lackluster. New architecture designs and implementations always come with an increased cost to manufacture, but this was hard to swallow, especially with no titles upon which to test those features.

But all is not lost. Competition drives down prices, motivates innovation, and brings a better solution for all of us. It's all down to AMD. Navi is the name of the game, and we should be expecting this some time in the second half of 2019. This 7nm GPU will be manufactured by TSMC, the same company pumping out AMD's Zen 2 chips. We don't know a huge amount about Navi just yet, other than rumors circulating that it's being positioned as a mid-range GPU right now, as opposed to anything capable of taking on the high end. This is interesting, because if it does pack a big enough punch at the right price in the mid-range, it may very well pull down the pricing of Nvidia's other offerings at the top, as it reshuffles its entire arsenal to compensate. And as we know from experience, it's the mid-range cards that hold the most clout when it comes to profit and market share.

That said, even if AMD doesn't make a huge amount of bank in the world of the PC enthusiast, its Radeon brand is assuredly safe, thanks to its immense dominance in the world of console computing, and, boy, is that a big market to control. Does that mean this is the end of AMD at the high end? Probably. Can we expect Nvidia prices to keep climbing? Well, unless another manufacturer comes out with a complete curve ball in terms of how we render in-game graphics, it's looking likely that its dominance in this market is assured.

Could we see an RTX 1060 without the RTX heavyweight hardware?

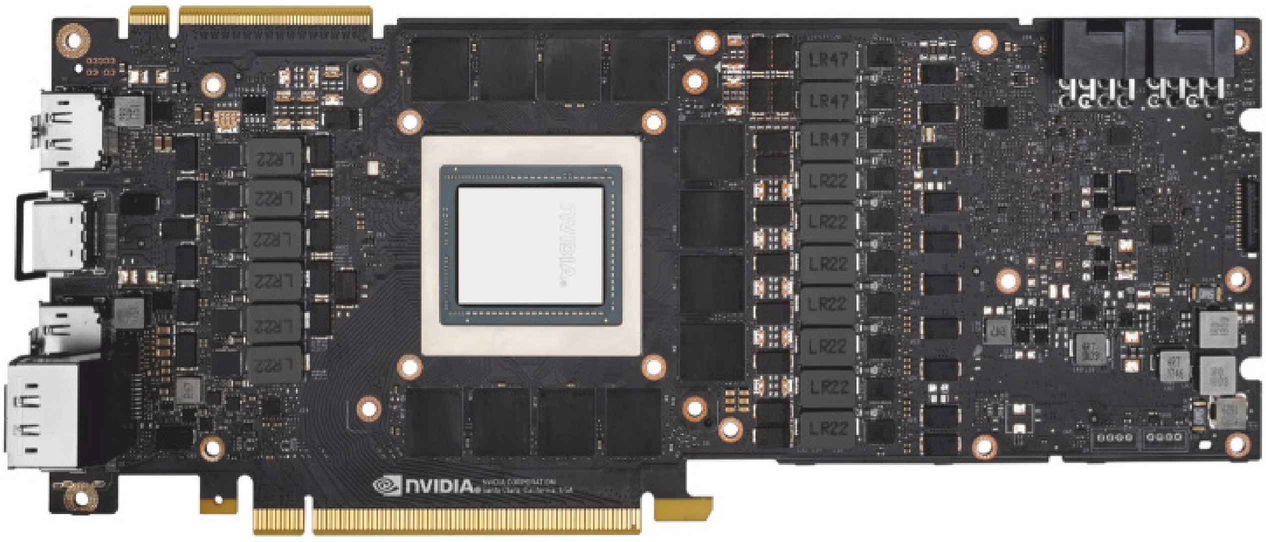

NVIDIA'S MID-RANGE STAKE

That said, we're still expecting a few releases from Nvidia some time this year: a Turing (or rebadged Pascal) variant of the RTX 2060, and we're seemingly still missing a Titan of some description, too, although we'd be surprised if that wasn't any less than $1,900, for one extra GB of GDDR6, a few CUDA cores, and some developer features. The 2060 is also of particular interest, and it'll mostly come down to the nomenclature Nvidia decides to use. If it's a GTX card, and Nvidia forsakes the RTX hardware (which would make the most sense, given its lackluster performance, even at 1080p, on the 2070), then it could be quite the impressive value card.

Spec-wise, we'd expect around 1,920 CUDA cores, and potentially either 8GB or 6GB of GDDR5X, as opposed to GDDR6. That should then perform quite nicely at around 60fps at 1440p. Although this is, of course, all speculative.

We've already seen some preliminary figures leaked online, most notably in *Final Fantasy XV's* benchmark, where a supposed RTX 2060 scored 2,589 points at 4K; in comparison, the GTX 1060 scored 1,985 points total. As expected, that's an increase of around 30 percent in performance.

Although this hasn't been confirmed officially, the card also retained its RTX nomenclature. If we did see RTX hardware make the cut, it'd likely be 24 RT cores and 192 Tensor cores—a little low for true ray-tracing or DLSS support. DLSS is arguably going to be more important at 1080p gaming than ray-tracing, but even so, there's still no word on price.

STORAGE SOLUTIONS SOLIDIFIED

Intel and Samsung duke it out

THIS TIME LAST YEAR, we predicted the future of non-volatile memory, and therefore SSDs, would be all about Intel's 3D Xpoint and Samsung's Z-NAND technologies. Fast-forward 12 months, and not much has changed, including the fact that we're still largely talking about the future. Neither 3D Xpoint nor Z-NAND have yet made much impact in consumer-level PCs.

That said, Intel's 3D Xpoint tech has become more widely available. Multiple generations and formats of drives based on 3D Xpoint and sold under the Optane brand have been released. We've had first-gen PCIe boards and M.2 cache cards replaced by second-gen products, while larger M.2 SSDs with just about enough capacity to run as primary drives have also been released. You could even argue that the mighty Intel 905p PCIe board is a third-gen product.

While Intel has made much progress, including refining its Optane controllers and algorithms, and bringing prices down, a limitation remains: Intel is still using first-gen 3D Xpoint memory chips. What's more, the Intel and Micron venture that created 3D Xpoint is on borrowed time. The two will part ways toward the end of next year.

The companies say they will keep working together to release second-gen 3D Xpoint memory chips in the first half of 2019, following which their venture will be done. For now, almost nothing is known about that second-gen 3D Xpoint tech. But more capacity at a lower cost would address the main remaining reservation we have with Optane drives: They're so expensive.

As for Samsung's Z-NAND, we've yet to see one in our labs, and no consumer variants exist. Despite that, Samsung has announced a second generation of the tech. It includes new SLC Z-NAND that doubles the data density of its first Z-NAND chips to 128Gb, and a new MLC variant with a nifty 256Gb per chip. However, the first products based on this will still be pricey, with large drives aimed at enterprise applications, as opposed to everyday consumers, a market shored up by its 3D NAND applications, and 970 Pro and Evo drives.

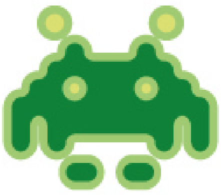

What's also unclear is the extent to which Samsung's Z-NAND is protected by patents. Toshiba's upcoming XL-Flash promises a very similar set of advantages as Z-NAND, including 10 times lower read latency than NAND. The overarching point involves competition. While Z-NAND and XL-Flash don't offer the same latency advantages as 3D Xpoint, they do promise to keep Intel on its toes.

How all that maps to SSDs isn't entirely clear. However, SSD prices are at an all-time low going in to 2019, and with QLC or quad-level NAND flash ramping up, large drives are only going to get cheaper. Will we see the first sub-$100 1TB SSD from a big brand in 2019? Maybe. At the same time, second-gen 3D Xpoint memory chips, plus wider adoption of Samsung's Z-NAND and perhaps Toshiba's XL-Flash, should enable a new generation of lower cost ultra-high performance SSDs. They still won't be cheap, but it's not hard to imagine a setup with something like a 250GB Optane OS drive and a 1TB SSD for mass storage being affordable by the end of the year. For most PC enthusiasts, however, 2019 will be about conventional NAND SSDs becoming cheaper than ever.

THE FUTURE OF DRAM

This could be the year we see the first DDR5 memory kits make it to market. We saw Micron and Cadence demo working DDR5 memory in late 2018, and even with JEDEC not finalizing the spec just yet, manufacturers are expecting to begin production in the latter half of 2019.

When it comes to DDR5, although it looks to be more of the same, it couldn't be further from the truth. Expect MT/s (MHz for the marketing gurus out there) to start at speeds of 4,266MT/s and reach 6,400MT/s, followed by a drop in voltage from 1.2V down to 1.1V, and a doubling in maximum capacity per chip, as usual.

According to Cadence and Micron, however, all that is just icing on the cake. The big changes will be at an architectural design level. Thanks to an improved command bus efficiency, superior refresh structures, and more bank groups being available to the processor, end users can expect nearly a 30 percent performance increase in real-world bandwidth, and that's just when comparing two identically clocked DDR4 and DDR5 kits together. Take a 3,200MT/s kit and compare it to a DDR5 4,800MT/s kit, and you could be looking at nearly 80 percent overall bandwidth improvements.

This is likely to hit servers at the end of 2019, with desktop applications ramping up mid-to-late 2020. It's possible, given AMD's release schedule, that we could see this launch with the Zen 3 architecture, as it will require a new socket, and as Zen 2's EPYC processors will be backward compatible with current sockets, it's unlikely Threadripper 3 and Ryzen 3 will require new sockets as well.

INTERFACE CONNECTION ADVANCEMENTS

Where will the world of interconnectivity take us?

RUMORS ABOUND that PCIe 4.0 is coming to the masses. Perhaps not in the early part of the new year, but you can bet it'll be here by the end of 2019. Announced in 2016, and ratified by PCI-SIG (a super-conglomerate of over 900 companies, including Intel, AMD, Dell, HP, and IBM) in the summer of 2017, PCIe 4.0 should have already been in production and on mainstream boards sometime in 2018. However, with a lack of support on the processor side of things, and there seemingly not being a need for the connection standard, delayed implementation and progress on introducing the spec has stagnated. We just didn't need it. Until now.

If Threadripper and Ryzen 3 follow the same or a similar specification as that of EPYC 2, both processor lineups should see support for PCIe 4.0 included as standard, with Intel also likely to follow suit with its 10nm parts. And with it being backward compatible, not including it on the next generation of motherboards would simply be ridiculous.

What's so great about PCIe 4.0? For starters, it doubles the per-pin raw bitrate, increasing from 8GT/s to 16GT/s. It also comes packing a reduction in overall system latency, along with improved flexibility and lane-width configurations for developers who are interested in tapping it up for lower-power applications and devices.

The big one for us, however, is that increased data rate. 16GT/s per lane should enable PCIe x4 devices to operate at up to speeds of about 6.4GB/s, versus the 3.2GB/s cap we're at now, giving us a lot more headroom for flash storage devices, and hopefully improving CPU-GPU communication and performance in multi-card situations, especially at higher resolutions. The latter is less of a problem for Nvidia, with its NVLink, of course, but for AMD, it's a pretty big deal, especially if the company gets back into the high-end GPU game with Navi.

PCI-SIG has also teased us with PCIe 5.0. Although complete specifications are sparse right now, we do know that it will be another doubling of the bandwidth limitations, meaning 32GT/s per lane, or 12.8GB/s theoretical maximum for transfer speeds on flash devices, and a staggering 128GB/s bandwidth for something like a 16-slot GPU.

The big question is how will these standards be introduced—if at all? Opinion in the office is split. The first camp, rooting for PCIe 4.0, reckons that thanks to EPYC 2's integration of PCIe 4.0, it wouldn't make a ton of sense to develop and assemble two separate processor designs, when you could produce one and disable some of the cores. Thus, as PCIe 4.0 is backward compatible with the other generations, adding it is a fairly easy task.

The other camp, the skippers, reckons that the bandwidth provided by PCIe 4.0 simply isn't enough, and because it takes both time and money to validate a spec like this (bear in mind that PCIe 5.0 has already been announced, and is due for launch in 2020), why bother, when we'll just need to do the same again in 2021? It's a tricky one, and only time will tell.

NVLINK INNOVATIONS

AMD might well have finished with high-bandwidth multi-GPU

SIX OR EIGHT CHANNELS?

DDR5 might not be on the cards until 2020, so is there anything we could see in the meantime to bolster memory performance? The most obvious solution would be an increase in the number of memory channels accessible by high-end desktop users, as that's typically where memory matters most. Mainstream platforms can get away with dual-channel and DDR4 for some time, but for the high-end, any increase in memory bandwidth bolsters performance in ray tracing, video rendering, CAD CAM, you name it.

AMD is likely to be unwilling to introduce more channels on its high-end motherboards just yet, because apart from including support in its latest Threadripper processor designs, it would also need to redesign its LGA socket on motherboards, so the newer chips would cease to be backward compatible with current X399 mobos.

That said, there is a small chance that AMD could pull the trigger on that sooner than we think, depending on what Team Blue does. We've known since August that Intel intends to launch a 28-core unlocked Xeon on its high-end consumer platforms soon. This 14nm 28-core, 56-threaded beast comes on a new chipset (X599), with a new socket (LGA 3647). It requires a boatload of power, and one hell of a VRM solution to compensate, but more importantly (because of its Xeon heritage), it includes a six-channel memory solution. There's proof of this already in Asus's ROG Dominus, as it packs 12 DDR4 DIMM slots, alongside two custom M.2 DDR3-style DIMMs. On top of that, the board requires four eight-pin EPS power connectors, dual six-pin 12V connectors, and dual 24-pin ATX mobo power. And did we mention it has an actively cooled 32-phase VRM design as well? Crazy.

That six-channel memory support is the first of its kind, and if the board and chip sell well before AMD's Threadripper 3 series arrives, the Threadripper 4 refresh may bring with it the eight-channel memory support found on its EPYC and EPYC 2 server-grade brothers.

What does it mean for end users? Well, if we take a dual-channel memory kit at 3,200MT/s, the total bandwidth is 51.2GB/s, quad-channel doubles that to 102.4GB/s, six-channel gives 153.6GB/s, and AMD's eight-channel memory would give 204.8GB/s, ideal for transferring large data sets or raw footage.

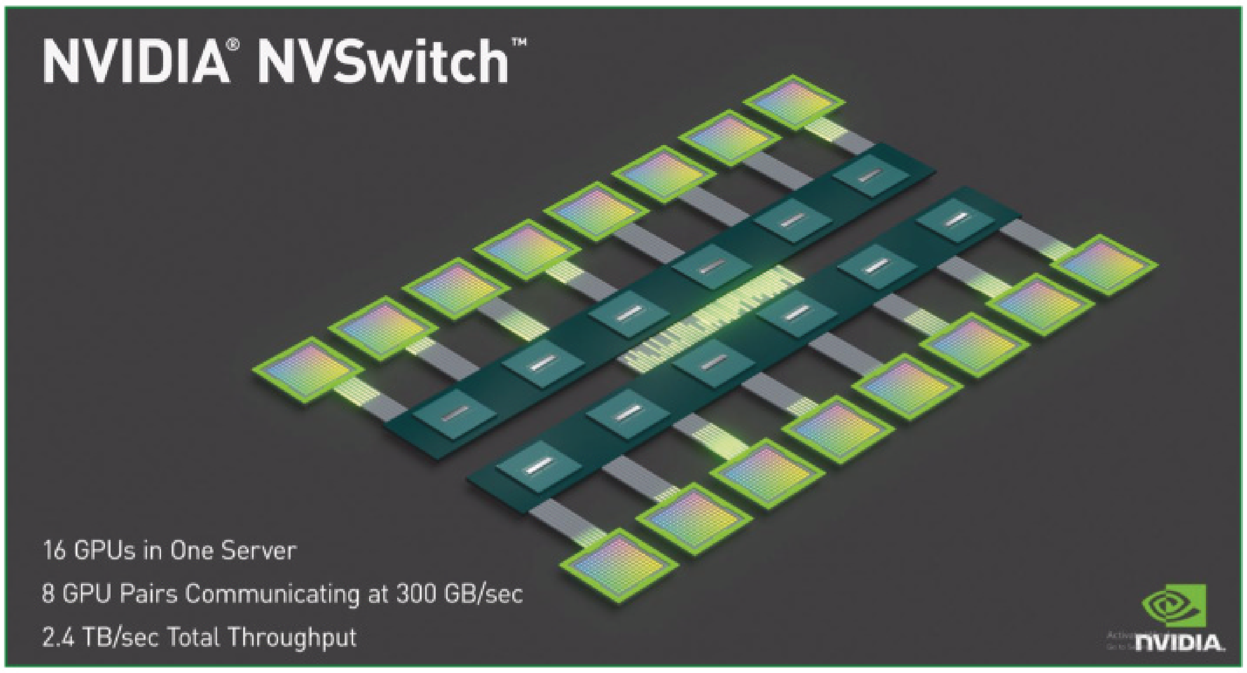
NVIDIA® NVSwitch™
16 GPUs in One Server
8 GPU Pairs Communicating at 300 GB/sec
2.4 TB/sec Total Throughput
NVIDIA

interconnects, but because the company isn't currently inhabiting the high end of the graphics world, it's easy enough to understand why it would think that they are unnecessary. Nvidia, on the other hand... Well, that's a whole different ball game.

Although the RTX series has come with many caveats, when it comes to both price and support for as-yet unreleased and unproven game development features, one thing that is true is the ridiculous amount of data those cards can pump out at any one time. The RTX 2080 Ti, in particular, can push over 150fps in some 1080p titles alone at max settings.

When running multiples of these cards, particularly in more extreme workstation environments, having a way of transferring data across the cards simultaneously is pivotal when trying to maximize performance. And that's exactly where NVLink originated: the server environment. In short, NVLink is a proprietary standard developed by Nvidia, which debuted back with the launch of its Pascal architecture, although the variant we know today was first seen with the Tesla V100 GPUs found inside the DGX-1 server. Each NVLink connection supports up to 50GB/s transfer rate (25GB/s each way simultaneously), well over 10 times what a single PCIe 3.0 channel can transfer, allowing for fast and seamless data throughput between the cards. The V100 cards supported up to six NVLink connections for a maximum transfer rate of 300GB/s

When it comes to mainstream cards, the RTX 2080 Ti supports two NVLink connections for a total of 100GB/s, while the RTX 2080 supports only one, for 50GB/s. Unlike SLI, NVLink doesn't act as a display interface, merging video outputs together like a scheduler, then back through the GPU-out. Instead, thanks to its super-low latency and high bandwidth, it allows multiple GPUs to pool their resources together as a shared entity, and communicate across the bridge, reducing traffic through the PCIe bus below, making the cards more efficient.

It's difficult to say just how much of an impact this will have on the gaming scene right now, but Hothardware.com suggests that in *Shadow of War*, at 1440p, there's as much as a doubling of performance when using NVLink versus traditional SLI, so far as scaling is concerned. ⏻

NVLink surpasses the world of the PCIe barrier.

PCIe 3.0 may finally be on its way out.

GIGABYTE
GIGABYTE B360M D3H
Designed in Taipei
Intel Optane Memory Ready
UEFI DualBIOS
Ultra Durable Intel GbE LAN
DDR4 Socket 1151

USB 3.2 AND THE TYPE C STANDARD

This will likely come as no surprise, but 2019 should see the advent of a lot more USB C-compatible devices, from monitors to GPUs, mass storage solutions, and eventually even peripherals. USB Type C will be the dominating form factor going forward, mostly due to its spec and form factor. Reversible by design, and secure-fitting, USB C cables are much easier to use than their USB Type A counterparts. Couple that with it supporting MHL, Thunderbolt 3, HDMI 1.4b, and DisplayPort 1.4 standards as well, and it starts to look like a winner. The biggy, however, is going to come with the introduction of the USB 3.2 specification. Ratified in early 2017, and still retaining its backward compatibility, USB 3.2 should see a doubling in total bandwidth, from around 10Gbps to 20Gbps (or 2.5GB/s).

QUANTUM COMPUTING

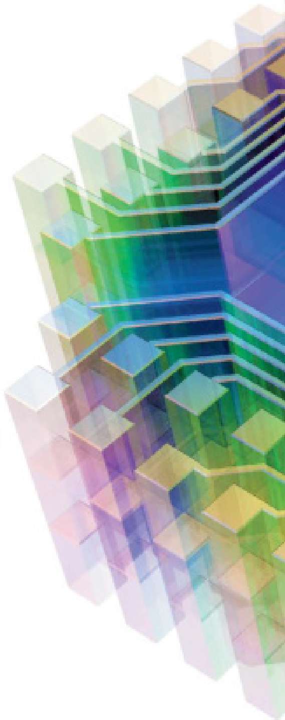

Mats Tage Axelsson introduces you to quantum computing, the coolest tech around; learn how it works, and how you can get started

QUANTUM COMPUTING has caught the attention of large companies, academics, and hobbyists. We're going to cover the history, the different ways to make a quantum computer, and the logic behind programming. You'll also learn about some programming toolkits that you can use to get started.

To run a quantum computer, the physics has to be understood, so programmers can manipulate and measure the final results. Scientists have observed quantum effects in photons, electrons, and isotopes of many materials. This means engineers use superconducting materials, such as niobium and aluminum, to construct workable quantum computing systems.

The logic gates are made of silicon wafers, and are controlled using microwave emitters. These solutions may not be the best in the long run, but they're the ones that are running now. To use quantum computers, you need logic that takes advantage of the two core concepts: superposition and entanglement. When you start exploring these concepts, the Bloch sphere helps visualize what to do with different gates. Programmers can use classical bit gates together with quantum gates to create the algorithms needed.

The mainstream media hails quantum computers as significantly faster than current models. It turns out they're only fast in specific areas: cryptography, optimization, simulation, and database searches. In cryptography, many algorithms are safe, because factorizing large prime numbers with classical computers takes far too long to be practical. Shor's quantum algorithm can do it in minutes. Optimization and simulations can benefit from a quantum computer's ability to test many solutions at once.

Database searches are faster by a factor of four. And with speedier database searches, machine learning also becomes much faster.

The history of quantum computers begins with quantum physics. In 1900, Max Planck first proposed

ENCRYPTION MAY BE AT RISK

One big question is whether encryption systems are at risk with quantum computers. The major concern is the RSA encryption scheme. The scheme is secure because it relies on the condition that factoring a large number into its primes is too time-consuming. When trying to find the prime numbers, there are many strategies, so the simplest one is to guess and try. A trial and error approach isn't practical, though, since a 2,048-bit number has millions of solutions. Some strategies can reduce the number of possible solutions, but even the most powerful methods will take years or millions of years. With the right algorithms, a quantum computer could reduce that time to a practical level. Efforts are underway to create other algorithms that are not breakable this way.

While this is prudent, the risk that a quantum computer can do this within 15 years is low. Today's quantum computers are both small and hard to program. The frameworks available for programming are few and far between. As you can see in other parts of this article, you'll still be setting a few qubit states and twisting the states. Converting that to a fine-tuned encryption-cracker is, more than likely, a far-off prospect for now.

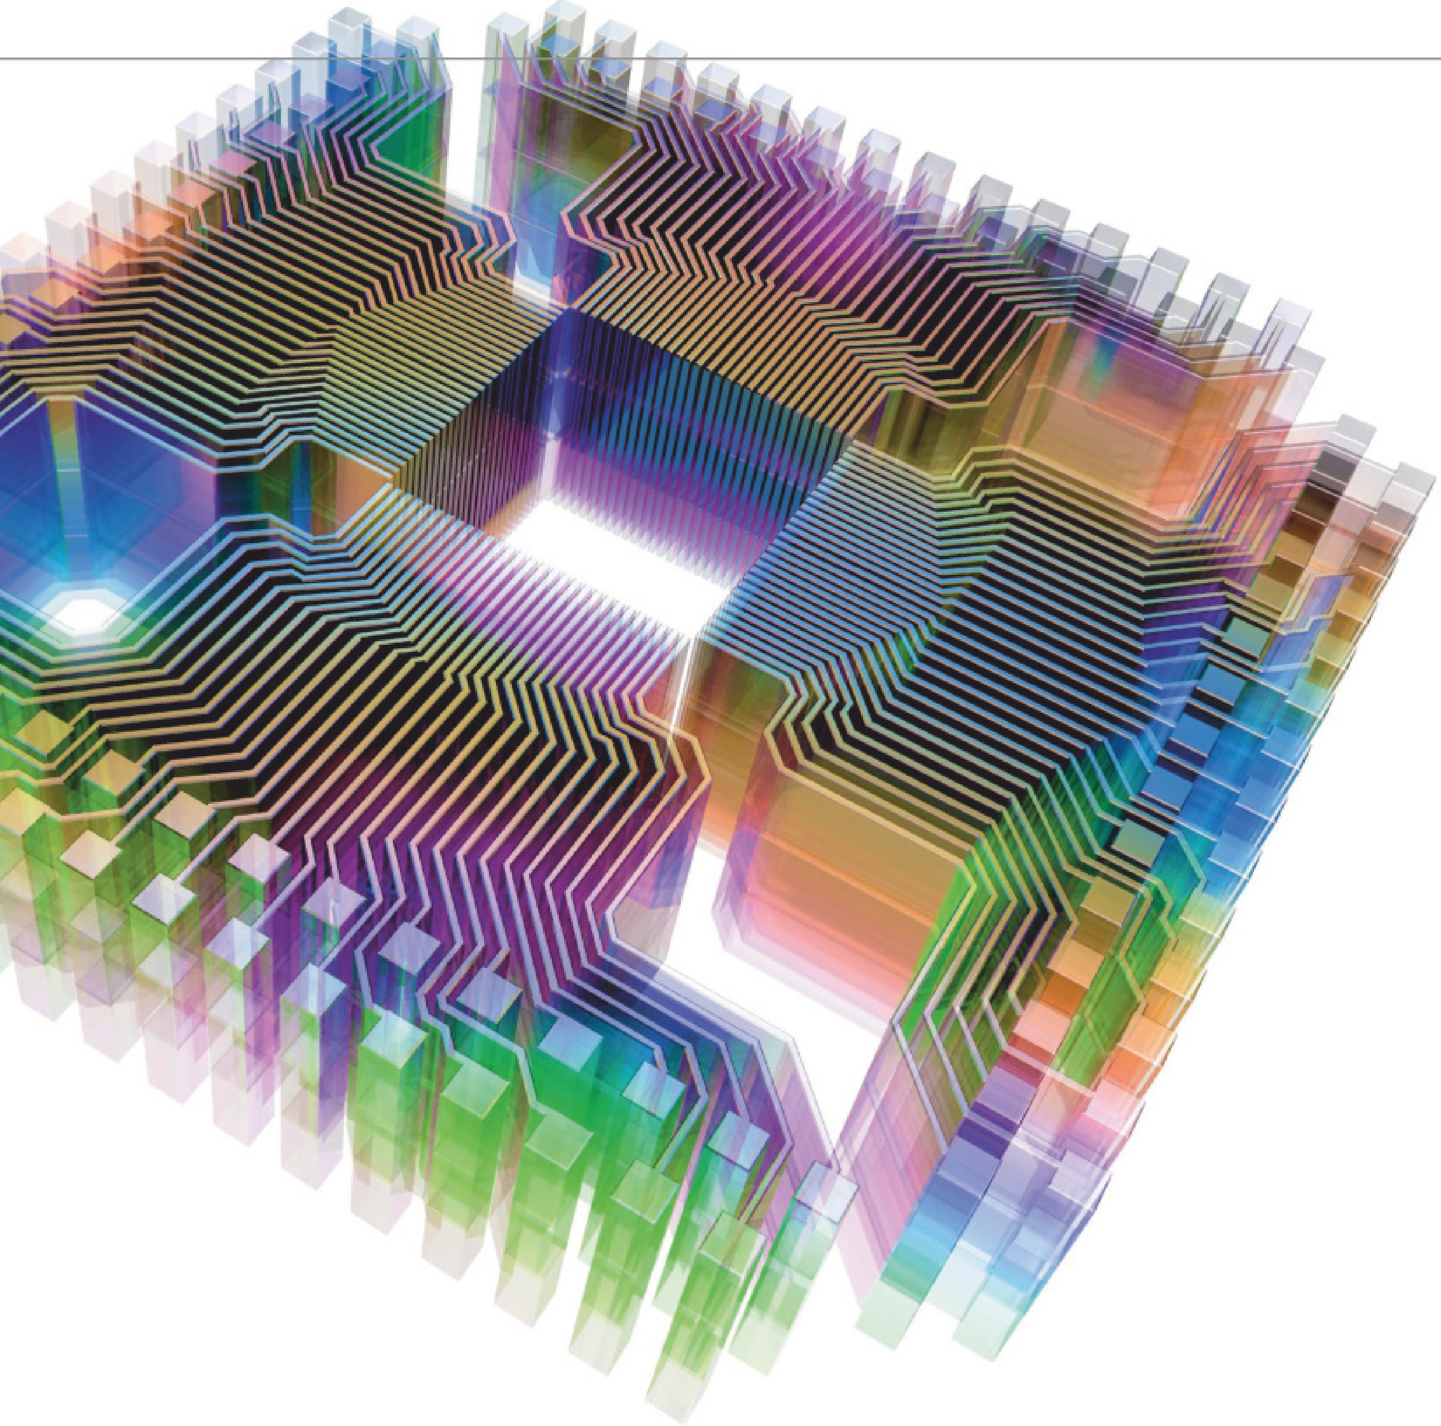

that light comes in discrete packets, called quanta. This discovery later led Einstein to show that Planck was right. When measuring the photoelectric effect, scientists could observe packet behavior. These two discoveries later led to quantum physics. Yet the deeper scientists dived into the quantum nature of things, the harder it got to explain how it works. For quantum computing, the most interesting developments are entanglement and superposition.

Superposition is where a particle exists in many positions at the same time. Considering this, scientists concluded that a quantum computer should be possible to build. A quantum computer is one that can do many calculations per operation.

THE CAT IN THE BOX

The famous physicist Schrödinger created a thought experiment. In it, he describes a cat and a gas container inside a box. The gas will poison the cat after the radioactive decay of an atom that releases the gas. The process is random, so observers won't know if the cat is dead or alive until they open the box. In quantum physics, this means that the cat is both alive and dead before anyone observes it. In the case of the cat, this is absurd, but in quantum physics, it's normal. What Schrödinger was saying was that there must be another explanation—one as yet undiscovered.

Entanglement is where two particles can have intertwined states. This means that the state of one particle is always the opposite of the other, even when they are apart. In quantum computers, the software creates entanglement; a CNOT gate creates this state.

To make use of all the phenomena demonstrated in quantum physics, scientists needed a way to describe what happens on a small scale. To program a computer, it needs logical operations, or logical gates. Quantum gates and classical logic gates are the same only up to a point. Quantum gates add features for changing states and entanglement.

It took until the 1970s for the first attempt at a theory to use quantum effects for computers. Shannon information theory describes classical gates and other aspects of data processing. For quantum computers, this is insufficient, because it doesn't specifically describe quantum effects.

Quantum information theory was first tried in 1976. During the 1980s, scientists made more progress, in part thanks to quantum computing conferences organized by MIT and IBM. Other developments included quantum cryptography and the first universal quantum computer.

To use all the states of the particle you measure, programmers need a formal language. Quantum information theory needed to improve. The different gates are the foundation for such a scheme.

Keeping track of what the different quantum states are is

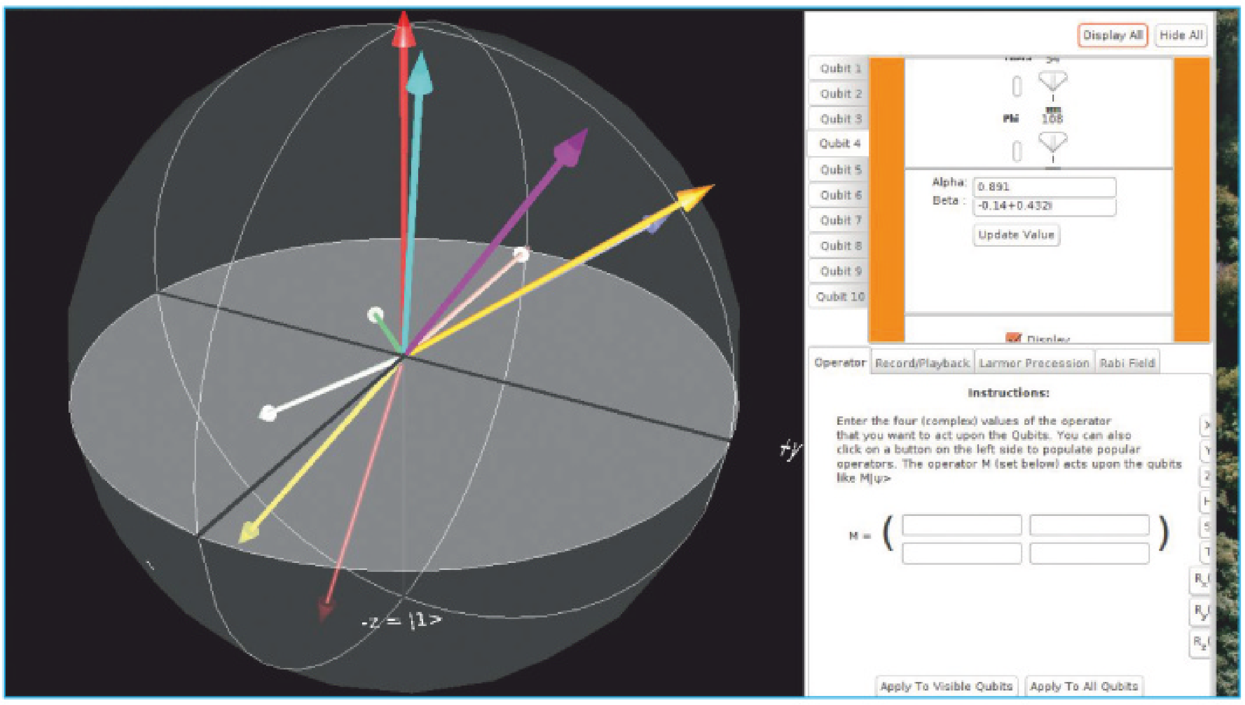

Alpha: 0.891
Beta : -0.14+0.432i

Instructions:

Enter the four (complex) values of the operator that you want to act upon the Qubits. You can also click on a button on the left side to populate popular operators. The operator M (set below) acts upon the qubits like M|ψ>

M =

Programmers use the Bloch sphere to illustrate how quantum gates manipulate the qubits.

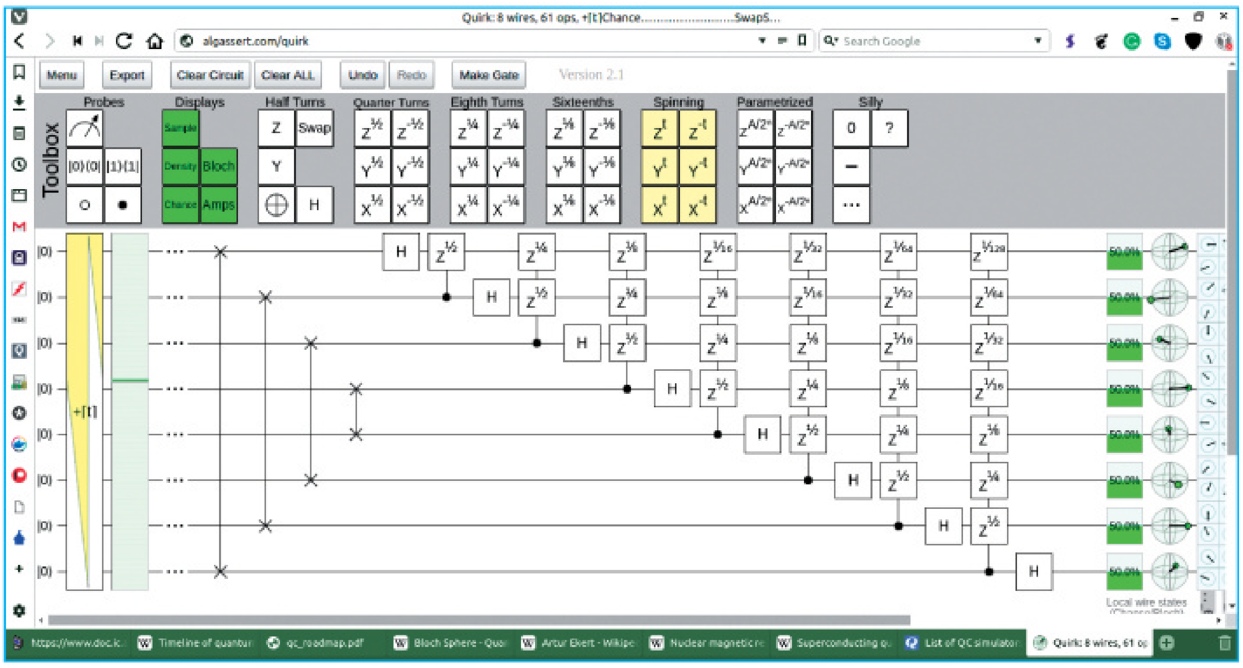

When starting out, programmers use a score, like a music score. The picture shows Quirk, an online simulator. The circuit is showing a Quantum Fourier transform.

confusing. First of all, the way it works is counter-intuitive at best. Second, there are many different spin axles to keep track of. The system isn't complicated, but it involves atypical approaches that are tricky to grasp. To make sense of all the state transitions, physicists created the Bloch sphere.

Peter Shor discovered Shor's algorithm in 1994. This algorithm solves the problem of finding the prime factors of large numbers. The basis of all encryption is that you can't, in a reasonable time, solve this problem. Quantum computers may solve this problem in minutes. Interest in quantum computers subsequently sky-rocketed.

Shortly after Shor's discovery, Lov Grover invented the Grover algorithm. This is best known as the database search algorithm, but it's also useful for other tasks.

Through the coming decades, scientists invented new features. Quantum error correction and fault-tolerant quantum computers were among the major ones. The first demonstration of a quantum computer took place in 1998. After this, development accelerated.

THE ION AGE

To date, there is a range of quantum computers in use. Researchers are using ion traps, light-based, and microwave-controlled Bose-Einstein condensates. These types need different control mechanisms. The ion traps use an electromagnetic field to trap the ions. Lasers create the field, keep the ion trapped, and "pump" the state of the electron of the ion. The electron will emit light only when the state matches the laser's frequency.

This is the logical 1 state, while the logical 0 state is the opposite. This makes it simple to measure, but achieving precision is still a challenge. Manufacturing is also expensive using current technologies. Meanwhile, a Bose-Einstein condensate is a gas cooled to millikelvin temperatures. When the gas reaches this temperature, all electrons take the same quantum state. Uses for this solution are, so far, limited to simulators—any actual computers are pipe dreams.

The interesting types are the Josephson junction types; these actually are running in practical setups. IBM's versions are even available for anyone who registers on the Q Experience web page.

The D-wave machine uses Josephson junctions. These are superconducting niobium and aluminum-oxide gates. At temperatures near absolute zero, a sensitive magnetometer measures their magnetic properties. The values show the quantum state of the junction, making the circuit a qubit. The principle this computer uses is quantum annealing.

IBM uses Nuclear Magnetic Resonance (NMR). The qubit is still a Josephson gate. That solution uses microwave transmitters to interact with the qubit.

The difference between the two is that the IBM solution controls qubits to a higher degree. This has made it easier for Canadian company D-Wave to manufacture bigger chips. At the same time, this limits D-wave to fewer applications. For example, an annealing system can't run Shor's algorithm or Grover's. It's useful for optimization and machine learning—no wonder Google already has one of its own D-wave quantum computers.

MORE IS LESS WORK

As mentioned, the algorithms in quantum computers are different. To explain what the difference is, you must bear in mind the physics involved. You probably already know that a conventional computer uses gates. The system uses Boolean logic to create the algorithms for calculating everything. The most important aspect of this is that it's deterministic, so we only deal in certainties. The answer is either yes or no—at least on the circuit level.

In contrast, quantum computing involves all answers having a probability attached to them. No

QUANTUM SIMULATORS

To help understand how quantum computers work, there are several simulators. There's no way that you can simulate one efficiently, but if you want to understand the underlying principles, many of these are excellent. The first one worth exploring is a Bloch sphere simulator. Without having at least a rudimentary understanding of the Bloch sphere, you'll have trouble understanding what the gates do. Learn more about it at http://eecs.ceas.uc.edu/~cahaymm/blochsphere.

There are many more simulators—the majority are the results of PhD thesis work. For that reason, most of them are not active, only good starts. The few mentioned here are useful for self-education.

The JQuantum simulator looks and behaves like a quantum score. Using it is a little cumbersome, but if you're practicing this while reading a textbook, then it becomes usable and clear. It also supports your own functions.

This is useful when you follow a course and want to test it on your own. The only other way to do that is to use your IDE and write a quantum program yourself. If you have Python experience, this is doable, but be warned—the simulation won't look pretty.

TOOLKITS FROM THE BIG GUYS

To program yourself requires the standard called OpenQASM. This is the basis for all development kits.

Many groups have developed toolkits using this standard. The best-known ones are from IBM, D-Wave, and Microsoft.

IBM decided to use Python to create Qiskit, and you can download this kit from Github. It also has many sources and demonstration collections available. You can learn all about the current development state from there.

In both Eclipse and Netbeans, all you need to do is import the code into a project and explore. Don't forget to install Python 3.5 or higher before you try to compile. IBM's QE has these examples in its Python toolkit. When you have the sources installed, you can only run simulations on your own computer. If you want to run on a real quantum computer, you need to get an account on the IBM Q Experience. The setup is simple: All you need to do is open an account and get your API token from your account. Then copy it into the "Qconfig.py" file of the project you're working on. There's a credit system for using the real machines, so simulate until you're certain that you've got it right. If you're really clever, you may be able to gain expertise level. In that case, you can obtain more units to run experiments.

If you're using Visual Studio, you can download the Qvis file and add the extension. You still need Python support, though.

answer is certain—rather, it is probabilistic. For this to be efficient, the algorithms must be adapted.

Quantum gates have an effect on one or several qubits at a time. When you measure qubits, the system usually does a complete analysis of the entire system. So, if you're searching for one qubit, that one will stick out like a sore thumb in one operation. Note that one operation is made up of many shots of the same gates. This is one of the reasons why quantum computers are faster than a conventional one.

The most common scenario is to find the special number of many, such as a database search. The other most cited case is for breaking encryption. Using Shor's algorithm, you measure the solution by the action of the gates. The gates run the input a fixed number of times before finding the answer. That number doesn't change when the input is a higher number. A classical computer forces you to test all possible answers. This means that when the number reaches 2,048 you are facing a potential million years to solve it. More numbers mean an exponential workload increase. The effort to solve the problem does not increase in a quantum computer.

LET'S HAVE A BLOCH PARTY

Once you know how many states and transitions you have in each qubit, it's time to code. Not so fast, though—there's a visual way to describe the gates. The Bloch sphere is the basis for all gates. The graph looks like a music score, so it's called a quantum score. It looks a bit different depending on the platform, but in general, the strings across are qubits. When you want to change the state of a qubit, you place a gate on the string. To create entanglement between two qubits, you drop the gate on one string and connect it to the other.

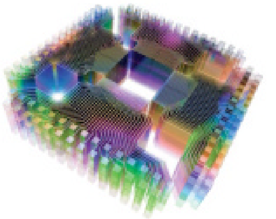

Here, we're employing Eclipse to develop our own quantum computing programs. Note that this particular toolkit is using Python as the supporting programming language.

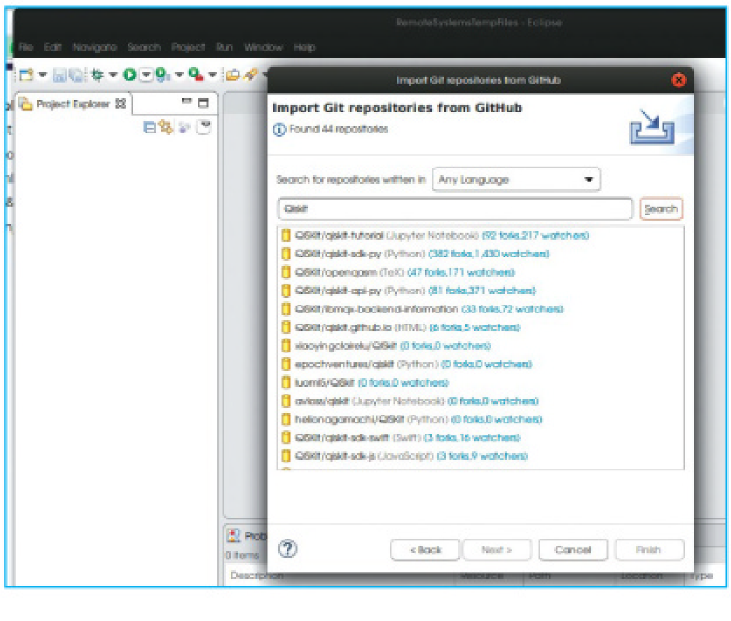
RemoteSystemsTempFiles - Eclipse
File Edit Navigate Search Project Run Window Help
Project Explorer
Import Git repositories from GitHub
Search
< Back
Next >
Cancel
Finish
Description
Resource
Path
Location
Type

Import Git repositories from GitHub

Found 44 repositories

Search for repositories written in: Any Language

Qiskit

- Qiskit/qiskit-tutorial (Jupyter Notebook) (92 forks,217 watchers)
- Qiskit/qiskit-sdk-py (Python) (382 forks,1,430 watchers)
- Qiskit/openqasm (TeX) (47 forks,171 watchers)
- Qiskit/qiskit-api-py (Python) (81 forks,371 watchers)
- Qiskit/ibmqx-backend-information (33 forks,72 watchers)
- Qiskit/qiskit.github.io (HTML) (6 forks,5 watchers)
- xiaoyingclairelu/Qiskit (0 forks,0 watchers)
- epochventures/qiskit (Python) (0 forks,0 watchers)
- luomi5/Qiskit (0 forks,0 watchers)
- avlass/qiskit (Jupyter Notebook) (0 forks,0 watchers)
- helonagamachi/Qiskit (Python) (0 forks,0 watchers)
- Qiskit/qiskit-sdk-swift (Swift) (3 forks,16 watchers)
- Qiskit/qiskit-sdk-js (JavaScript) (3 forks,9 watchers)

0 items

To describe the quantum state of a particle, you use spin, like a spinning ball. More precisely, the angle the "ball" spins. Even though there are no balls in the classical sense, the math describes it well. When the equipment measures an electron, it's in spin up or spin down in one axis. The system can only measure in one dimension at a time. The three are x, y, and z. If you measure x, you lose all information about the spin state on the z and y axes. When you describe the quantum state, you use an imaginary arrow in the Bloch sphere.

The top of the sphere represents the qubit's logical 0 and the bottom is logical 1. These are the only two states that the system can measure with 100 percent certainty. To get a result in a quantum computer, you make many "shots" and get the most probable answer.

A matrix contains co-ordinates that show the arrow's angle. The length of the arrow is always 1. The values we use to calculate the probability describe the components of the axes for the arrow. The Bloch sphere rules the gates and their naming.

The available gates depend on the application and type of quantum computer. But the following are the common ones. The Pauli X gate rotates the state around the x-axis, changing the value of z. This gate is also a bit-flip gate. The Pauli Y gate rotates around the y-axis, changing the x value and the z value. This gate is a bit-flip and a phase-flip gate.

So far, the gates only flip the bits and phases. These don't deal with superposition or entanglement. The Hadamard gate puts the qubit in a superposition. To make that clear, the qubit will have a 50 percent chance of being 0 or 1, flipping between X and Z.

It will act like a dice, only using two in a row will always give the starting value. More gates that extend the Hadamard gate are S and its transformed conjugate.

To use entanglement, you use a CNOT gate, usually depicted with a plus sign in a circle and a line down to another qubit. The gate flips the target qubit when the control is 1.

Those are the basic gates you'll encounter as you start out. On the IBM Q Experience, you have the chance to also create subroutines. A better idea, though, is to use a development kit. There are many available, and you can extend your IDE with them. Eclipse, Netbeans, and Visual Studio have modules. They use several languages. The programming tools for Python, for example, are available from IBM. Note that the routines prepare scripts for the quantum computer—it doesn't run Python on it.

So get cracking! You won't create much at first, but we guarantee you will have a lot of fun.

TRION 100
HyperX
SAVAGE
INTEL® SOLID-STATE DRIVE
750 SERIES
SAMSUNG
SSD 950

SOLID STATE OF THE NATION

With SSD prices at an all-time low, *Jeremy Laird* evaluates your solid-state storage options

ARE SSDS FINALLY cheap enough to qualify for mass storage? As we go to press this issue, that's surely an affirmative. Compared to the prices of the first mainstream consumer SSDs of a decade ago, you now get over 10 times the storage for about one quarter of the money. Back in 2008, an 80GB drive from a major brand would set you back around $600. Today? Prices for solid-state drives in the 240GB to 256GB class kick off around $35. For 480GB to 512GB models? Make that about $60.

Granted, you'd be brave to bag one of the cheaper options from lesser-known brands. But big-brand, full-featured 1TB SSDs can currently be had for as little as $140. For many PC users, that's both cheap enough and large enough to be considered viable for the mass storage of data. Of course, if you need major quantities of memory or simply want to absolutely maximize your storage space for as little money as humanly possible, traditional magnetic drives remain your weapon of choice. But a terabyte or three is enough for most of us, and in that context, solid state is now a serious option.

If that sounds like a simple sales pitch for solid state, the practical reality of choosing a drive is way more complicated. SSD technology has progressed beyond recognition since the first mainstream consumer drives hit the market 10 years ago or so. With that progression has come complication. Put simply, there's way more to a good SSD than impressive peak read and write numbers. In terms of the end user experience, other factors, including latency and random access, can be more important. Endurance and longevity—and how those issues map to both drive technology and usage scenarios—matter, too.

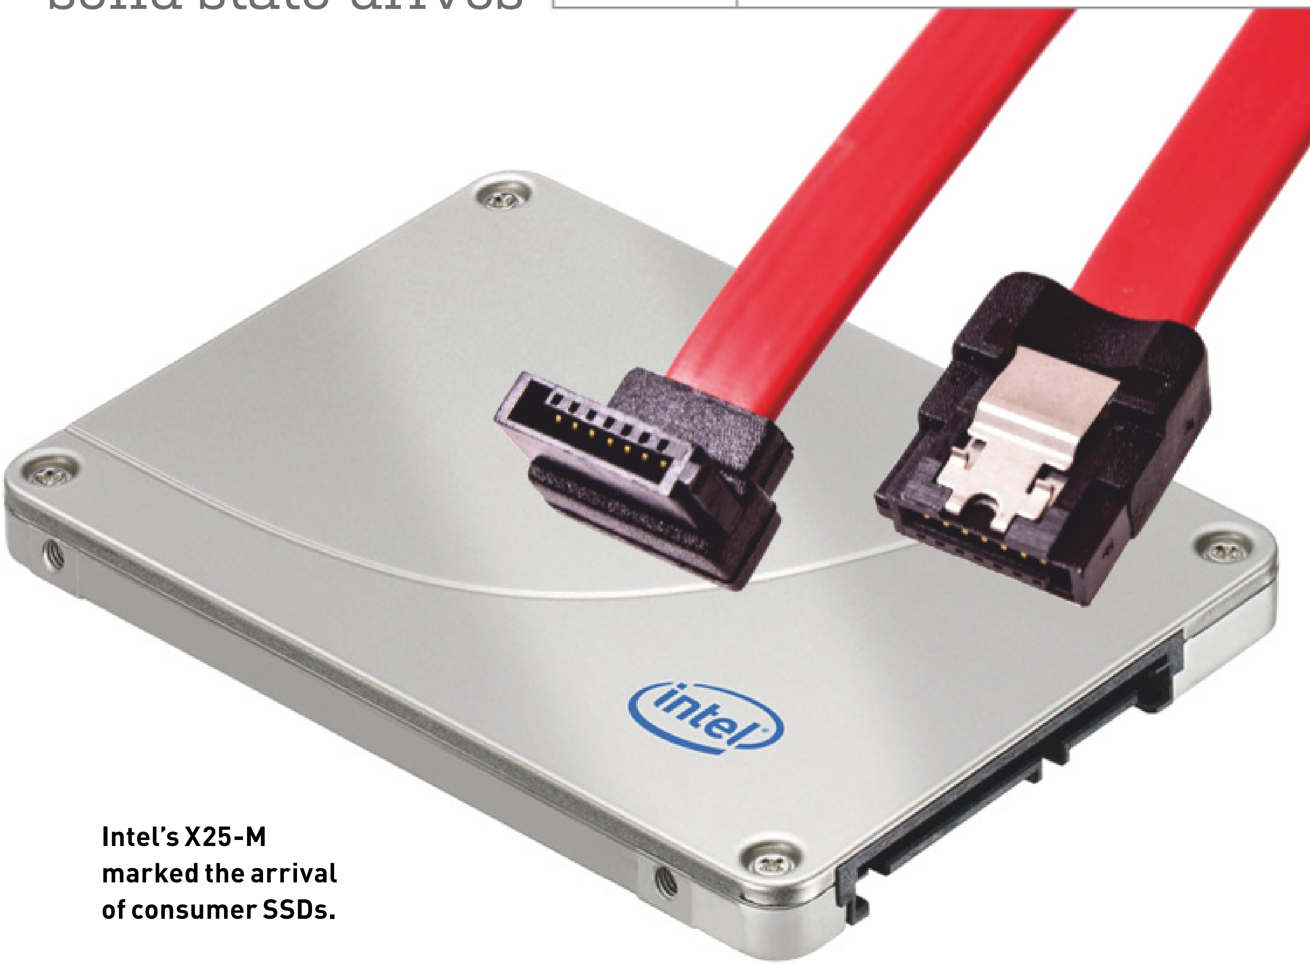

Intel's X25-M marked the arrival of consumer SSDs.

SATA limits SSD performance, but remains a popular choice.

Even at the component level, let alone full retail drives, there are far more vendors of SSD controllers and memory chips than there are makers of graphics chips and CPUs. It's tough to keep up with all the options. Likewise, working out how to actually judge the performance of an SSD is far more complicated than for other core components. But that's OK, because we're here to give you the low-down on the current solid state of the nation.

FLASH SALE

Wind the clock back to fall 2008, and the arrival of Intel's hot new X25-M solid-state drive. It wasn't the first SSD aimed squarely at the consumer market, but it did mark the arrival of SSDs based on NAND flash memory as a consumer technology, albeit at a price. The aforementioned 80GB drive for $600? Yup, that was an Intel X25-M.

Anyway, what matters is that the X25-M and its early-SSD siblings seemed like the final piece of the puzzle for achieving the holy grail that is the solid-state PC. Kiss goodbye to magnetic hard drives with spinning platters.

It's easy to forget just what a revolutionary leap it was to move from traditional hard drives to solid state. Moving parts, put simply, mean latency and unreliability. As it turned out, however, the step change to solid-state storage didn't go smoothly. After basically blowing everybody's socks off with HDD-battering headline performance numbers, it quickly became apparent that something was rotten in the solid state of Denmark. To be fair, the Intel X25-M wasn't the first problematic SSD, and it wouldn't be the last. But it was a very high-profile representative of the breed.

The problem is that the performance of SSDs configured with NAND flash memory degrades over time. The reason goes something like this: NAND flash memory is made up of cells. These cells are arranged in a hierarchical fashion. Individuals cells are arranged into strings and, in turn, arrays. These arrays make up pages, at which point the overall number of cells typically clocks in at 32,000–128,000 per page. Those pages are arranged into yet larger blocks, which ultimately measure in megabytes.

So what, you ask? Here comes the kicker. While NAND memory can be read and written in pages, it can only be erased in whole blocks. The reason for that is complex, and involves voltage levels and minimizing errors. But the impact on performance is sobering. Writing to empty memory cells is a breeze. But overwriting existing data is far more laborious. All data from the relevant block must be copied to a cache memory, the full block is erased, and then rewritten with the modified data.

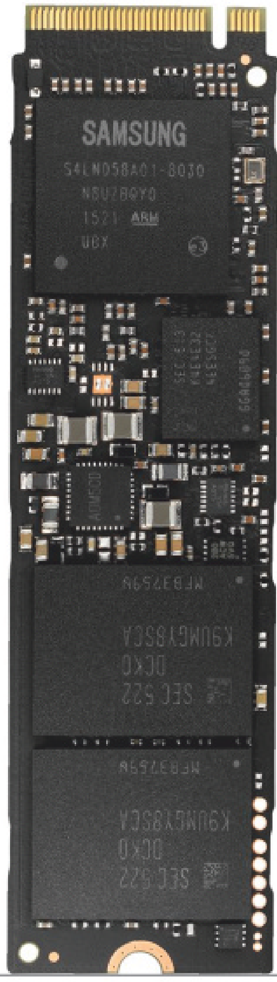

M.2 is now the default option for high-performance SSDs.

As a drive approaches full capacity, and the availability of empty cells, arrays, pages, and blocks dwindles, it's not hard to see how this impacts performance. But even when a drive has significant free capacity, performance can be dramatically compromised. Again, it comes down to the fact that only full blocks can be erased. The consequence is that, in the first instance, when pages within a block are deleted at operating system level, on the drive they're only marked as dead or redundant. To actually erase those pages would require erasing the whole block, and thus caching and rewriting any pages that retain live data. It's therefore expedient in the short term to simply mark pages as containing redundant data rather than erase them.

The long-term result is that, as an SSD fills up, juggling data becomes ever more laborious. A nearly full drive involves lots of long-winded read, cache, erase, modify, and write processes. But the same applies—or at least did apply—to a heavily used drive with lots of spare space. The upshot of it all back in 2008 involved very expensive SSDs that became stuttering, almost useless lumps of silicon.

Several mitigating technologies have since emerged, including intelligent garbage collection and the TRIM command, that largely solve the problem of extreme SSD performance degradation. But the saga remains significant in the

Samsung Evo 860 *(left)*; **Samsung Pro 860** *(right)*. **The SSD was the final piece of the solid-state PC puzzle.**

SAMSUNG
Solid State Drive

sense that it presaged a broader problem. Assessing SSD performance is much more complicated than merely measuring peak throughput from a box-fresh drive.

Hold that thought. We'll come back to the question of how SSD performance and indeed SSD benchmarking has developed since those early days. First, let's consider the current state of the SSD market, and have a look at all the technologies and options available.

The very basics of SSDs involve form factors and interfaces. Currently, two options dominate. The first comprises drives contained in 2.5-inch enclosures and hooked up via the legacy SATA interface, courtesy of cables. More recently, drives presented on 22mm-wide bare circuit boards, and which communicate predominantly over the PCI Express interface, have become popular, and are known as M.2.

SATA-based drives are most obviously limited in terms of bandwidth. The latest SATA III spec tops out at a practical maximum of 600MB/s. M.2 and indeed other PCI Express-powered drives are limited by the number of PCIe lanes they utilize. M.2 drives currently max out at four PCIe 3.0 lanes, and so 3.9GB/s of practical, usable bandwidth. Lesser-seen options include SSDs configured as pure PCIe cards and U.2 drives, the latter theoretically combining the best of both SATA and PCIe, but they haven't caught on in desktop systems.

Further nuance involves control protocols. Long story short, the legacy AHCI protocol used in conjunction with the SATA interface was never conceived

2.5-inch SOLID STATE DRIVE
crucial
by Micron

Crucial's MX500 is the obvious SATA option.

MORE OR LESS MAINSTREAM

When it comes to value SSDs, there's such a thing as false economy. A cheap drive from an unknown brand looks great on paper, with plenty of capacity and appealing performance numbers. But therein lies danger.

Engineering a flash-based SSD that delivers good long-term performance and reliability is no simple task. It requires deft integration of multiple components and quality coding of the drive's firmware. Specifically, one of the major challenges includes effective garbage collection that mitigates the inherent loss of performance over time as NAND blocks fill with data. For that you need a quality controller chip and good firmware algorithms. Bigger brands can invest in such technologies and share the cost over a wide range of products.

Another problem for cheap drives involves the related issues of over-provisioning and cache memory. Over-provisioning of memory allows a drive to maintain capacity even as wear causes cells to fail. It also allows a chunk of memory to operate in faster single-level cell mode. The very cheapest drives typically offer limited over-provisioning for wear and to act as SLC cache. While we're talking cache, cheap drives often also omit a DRAM cache, which will further impact their performance.

With all that in mind, what do we recommend? In the 2.5-inch SATA market, we would go for the Crucial MX500. It's a little more expensive than its BX500 sibling, but it's properly configured with both DRAM and dynamic SLC cache. Expect to pay about $50 for a 250GB drive and $80 for the 500GB version. 1TB is yours for around $150.

As for a mainstream M.2 NVMe option, we elect the Samsung PM961. It's a bare-bones OEM version of the popular 960 Evo, complete with the same TLC NAND, Polaris controller, and DDR3 DRAM. 256GB variants currently start at around $65. It's a lot of SSD for the money.

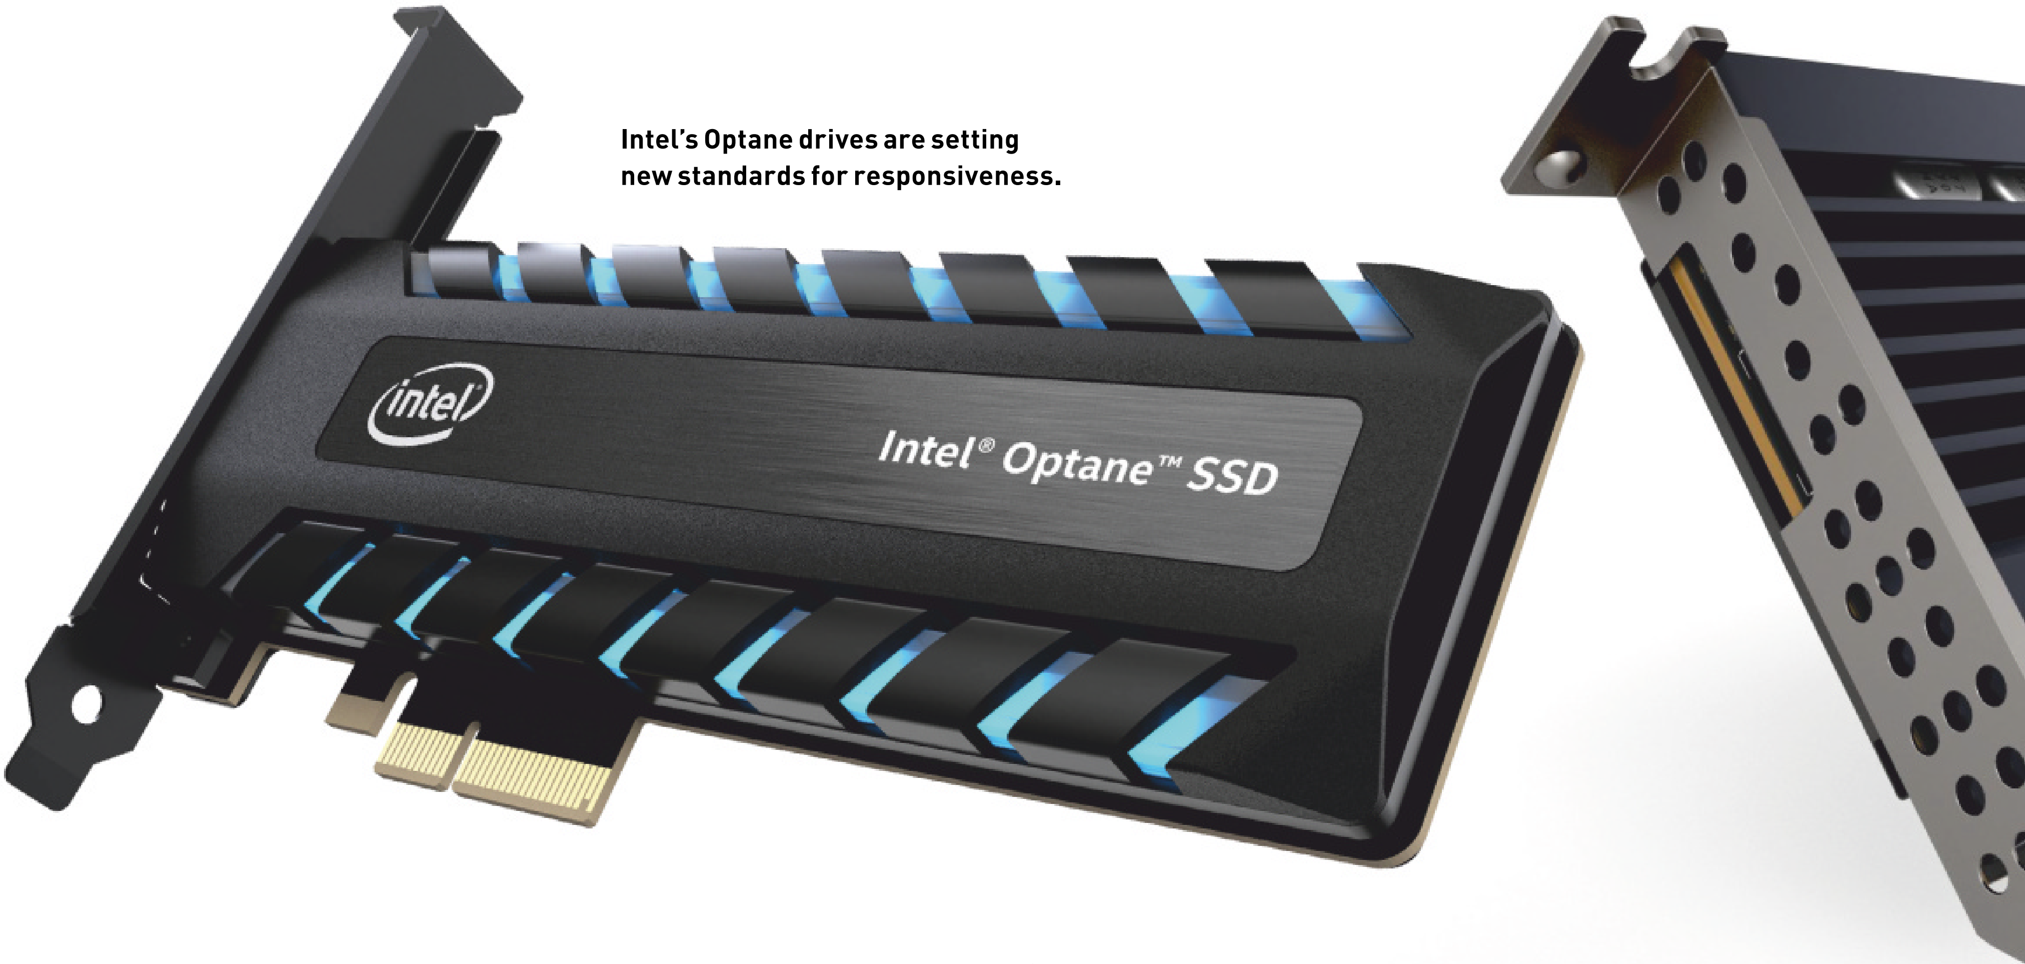

Intel's Optane drives are setting new standards for responsiveness.

with solid-state drives in mind. However, the newer NVMe (Non Volatile Memory Express) protocol solves that problem, and helps make the most of the theoretical advantages of solid-state tech, especially with random access performance. Most but not all M.2 drives support NVMe, so bear that in mind when shopping.

TAKE IT ON BOARD

Of course, motherboard support is central to much of this. Older boards do not have M.2 slots. The good news is that converter cards for M.2 SSDs that plug into standard PCI Express slots are available. But remember that legacy motherboards may not play nicely with newer SSDs, especially those that utilize the NVMe protocol. Our advice for any motherboard that does not have an M.2 slot is to check support for booting from NVMe drives before assuming it's present.

It's also worth considering platform support for PCIe-powered drives, generally. With Intel's mainstream and mobile platforms, M.2 slots are hooked up via the chipset, rather than directly to the CPU's own PCIe lanes. That means the drives ultimately communicate with the CPU via the DMI bus. As it happens, the current implementation of the DMI bus is effectively a quad-lane PCIe 4.0 connection by another name. Problem is, of course, the DMI bus carries all chipset bandwidth, including USB, legacy storage, networking traffic, and potentially multiple M.2 SSDs.

You can get around that limitation by using a pure PCI Express SSD or a PCI Express M.2 adapter card that drops into one of the available x16 PCI Express slots, and thus connects directly to the CPU. However, because Intel's mainstream desktop processors only have 16 available PCI Express lanes, using both PCIe16 slots means dropping any installed graphics card down from 16 to eight lanes, which could impact gaming performance in extreme circumstances. However, for both Intel's high-end desktop and all of AMD's existing desktop platforms, there are enough available direct-to-CPU lanes to maintain a full 16-lane graphics interface alongside a quad-lane PCI Express storage solution.

The final motherboard-related concern is to ensure you have an optimal configuration. When boards offer multiple M.2 slots, they may not be equal in terms of PCI Express lanes. What's more, not only will slots share the DMI link on a mainstream Intel board, as mentioned above, but they will conflict with the SATA storage bus. Put simply, with any given motherboard, there will be both limitations to and optimal configurations for connecting a mix of M.2 and SATA drives. Consult your motherboard manual for more information.

SSD Benchmark
Run
SCORE
SCORE
SCORE

Real-world SSD performance is about much more than raw throughput.

Next up is the actual technology used for the memory chips in an SSD. For better or worse, conventional NAND flash memory remains by far the dominant non-volatile memory technology in SSDs. Intel's Optane and Samsung's Z-SSD drives, it's true, offer potentially revolutionary non-volatile memory alternatives in the shape of 3D Xpoint and Z-NAND, respectively. However, it's early days for both technologies (see "The Future of SSDs" boxout on page 43).

But even among NAND SSDs, there is now a dizzying array of memory cell technologies. Just for starters, you've got SLC, MLC, TLC, and QLC NAND. Respectively, each technology can store one, two, three, and four bits of data per memory cell. The benefits of increasing per-cell capacity are obvious enough—more capacity for less die space, and so lower cost. But as cell density increases, there are trade-offs.

Firstly, endurance falls off a cliff. SLC or single-level cell NAND flash memory is good for as many as 100,000 program-erase cycles before a given memory cell loses its ability to store data. That falls to

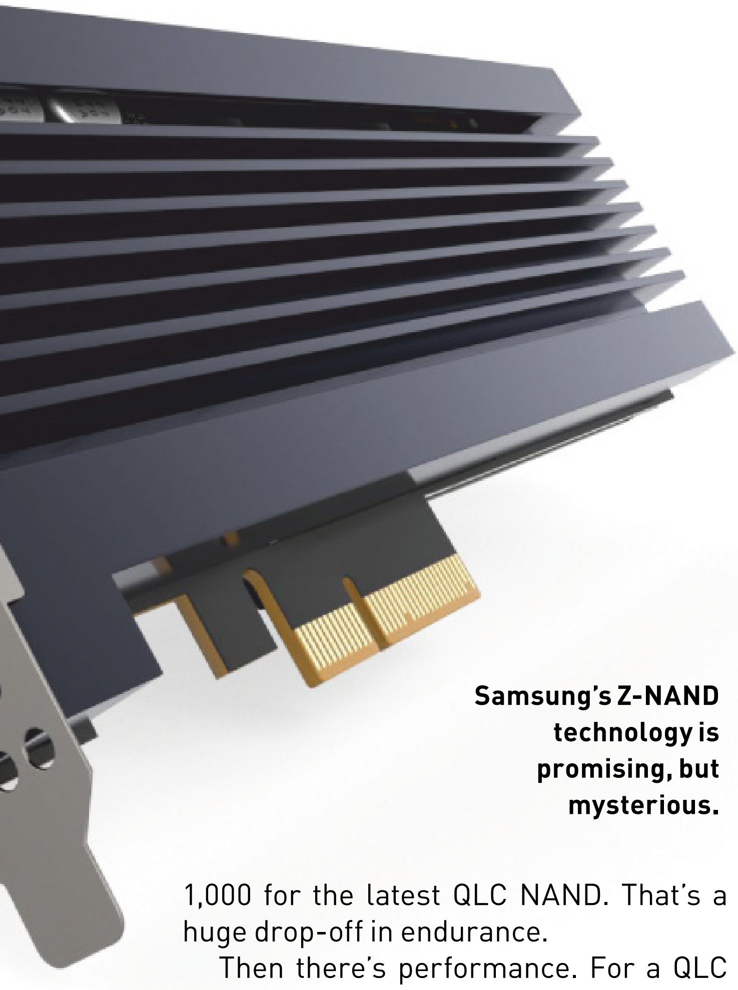

Samsung's Z-NAND technology is promising, but mysterious.

1,000 for the latest QLC NAND. That's a huge drop-off in endurance.

Then there's performance. For a QLC memory cell, reading and writing the contents of a cell means distinguishing between and applying 16 possible voltage levels. Long story short, reading and writing to multi-level NAND cells is an increasingly laborious process as per-cell density increases.

ALGORITHM SECTION

Inevitably, mitigating technologies that offset both the loss of endurance and performance as cell density increases have emerged. Wear-leveling algorithms ensure that writes are spread evenly across the drive, while over-provisioning of memory means any given drive can tolerate a certain amount of cell failure before available capacity begins to shrink. Such measures are good enough for Crucial to rate its new QLC-powered drive, the P1, with a mean time to failure of 1.8 million hours, and an overall write endurance of 200TB, which should be plenty for most consumers.

Performance, meanwhile, is much improved by allocating a large chunk of the memory as a cache running in much faster SLC mode. But as the drive fills up and eventually eats into that SLC cache, performance begins to tail off. It's another example of how SSD performance isn't entirely straightforward. It also helps to explain why big, cheap drives are inevitably slower.

As for 3D NAND memory, that's available in multiple per-cell densities just like standard 2D NAND, but involves

intel
Intel® Optane™ SSD 905P

Intel's Optane 905p is almost certainly the fastest SSD in the real world.

PREMIUM PERFORMANCE

When it comes to M.2 NVMe SSDs based on flash memory, we're not going to pull any big surprises. Samsung's 970 Pro, available in 512GB and 1TB formats for $190 and $380 respectively, is our pick of the premium M.2 drives. Its advantage over the competition isn't enormous, but it's still the daddy—just by a smaller margin.

Its core specifications are the chief reason why the 970 Pro is so good. For starters, it uses Samsung's hottest 64-layer 3D MLC NAND memory, where the 970 Evo uses cheaper but ultimately slower 3D TLC NAND. Then it packs at least 512MB of LPDDR4 DRAM cache. Finally, there's Samsung's very latest and greatest controller chip and firmware combo, known as Phoenix.

The net result? Sequential read speeds of 3.5GB/s that pretty much saturate the quad-lane M.2 interface. It's about as good as it gets for 4K random read performance at low queue depths, too. As good as it gets, that is, for a NAND flash drive. Unfortunately for Samsung, Intel has finally begun to deliver with its 3D Xpoint technology, now branded Optane. Which is why our money-no-object recommendation is the Intel Optane 905p 1.5TB.

It costs a preposterous $2,100. It doesn't even offer the highest peak sequential speeds. But it's almost certainly the fastest SSD you can buy in terms of real-world performance where it matters.

Most obviously, the 905p absolutely kicks butt when it comes to latency and low queue depth 4K random access performance. So, yes, the Samsung 970 Pro edges it by about 10 percent when it comes to sequential reads, but the 905p absolutely monsters the Samsung when it comes to single queue depth 4K reads—it's 300 percent faster. Combine that with the lowest latencies in the business, and you have the snappiest-feeling SSD on the planet. It's just a pity that it costs quite so much.

multiple layers of memory cells, each atop the others, instead of a simple planar 2D array of cells. Apart from the obvious benefits for memory density, 3D NAND also has inherent performance and endurance advantages. That's why Samsung's fastest NAND-based SSDs, including the 970 Pro, use 3D NAND.

Aside from interfaces, control protocols, and memory type, SSD controller chips are another major determining factor when it comes to performance. Unfortunately, relatively little is typically known about any given SSD controller, and it's usually impossible to make direct comparisons. But suffice to know that cheaper drives will usually make do with lower-specification controllers. Typically, such controllers will be capable of high theoretical peak speeds, but struggle when it comes to the finer detail of SSD performance. Quoted IOPS performance can be a useful yardstick when comparing drives in that context.

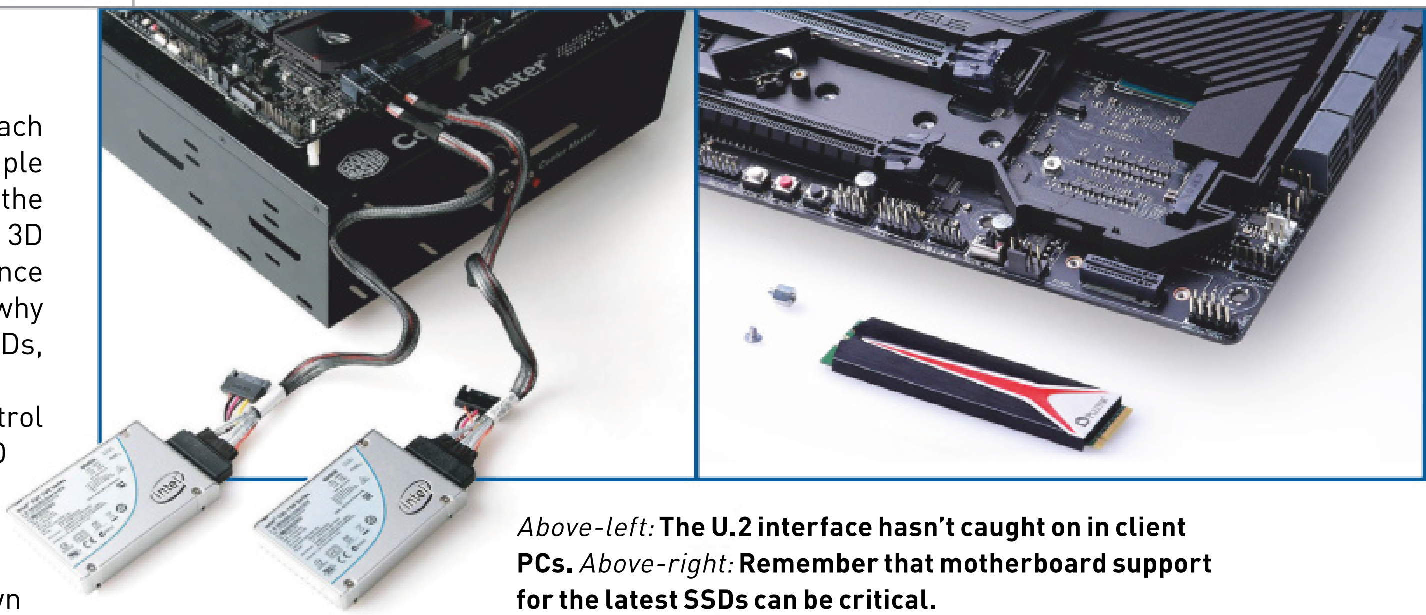
Above-left: **The U.2 interface hasn't caught on in client PCs.** *Above-right:* **Remember that motherboard support for the latest SSDs can be critical.**

Which brings us very neatly back to that question of SSD performance and, indeed, SSD benchmarking. It turns out that peak sequential performance isn't actually that critical to real-world PC performance. Instead, most storage traffic depends on rapid access to small chunks of data dotted around the drive. That's what makes for a fast and responsive-feeling desktop or laptop PC. The size of these small chunks is dictated by the smallest block sectors of storage, which for modern HDDs and for SSDs is 4KB. Hence the shorthand for this workload is 4K random access.

What's more, assessing 4K random performance also involves the question of queue depth. Put simply, queue depth is a measure of the number of simultaneous drive access requests. Because SSDs typically have multiple access channels, they can execute access requests in parallel, and serve up bags of bandwidth. That's quite common for a server PC delivering webpages or database access to scores of clients.

For desktop PCs? Not so much. Instead, it's very low queue depths of one to four that matter, and in those scenarios, an SSD's inherent parallelism counts for a lot less. Long story short, if you're looking for a single metric that gives the best insight into how snappy a given SSD will make your PC feel on a day-to-day basis, it's the 4K random read with a queue depth of one benchmark that you should seek out. Peak sequential performance simply isn't all that. ⏻

techradar.

THE FUTURE OF SSDS

Arguably, the future is already here. The two most promising candidates for technology to replace NAND flash as the go-to non-volatile memory technology are thus far Intel's 3D Xpoint (developed with Micron) and Samsung's Z-NAND.

Both are available today in real, buyable SSDs. However, neither has achieved anything like mainstream pricing and they remain largely targeted at high-end commercial applications.

3D Xpoint, which forms the basis of Intel's Optane drives, is a truly revolutionary departure from NAND flash. Instead of creating memory cells out of tiny capacitors as per NAND, 3D Xpoint is a resistance-based technology that uses bulk property change to the cell material to alter its resistance level and so record a memory state.

So, how does Samsung's Z-NAND compare? Like 3D Xpoint, the big selling point is low latency rather than a leap in sequential throughput. Samsung says its recently announced second-gen Z-NAND SSD, the SZ1735, offers five times lower latency than the best conventional NAND drives, and gets within spitting distance of 3D Xpoint. For comparison, Z-NAND is rated at 12–20µs (microseconds) for read latency, while 3D Xpoint achieves 10µs.

How does Z-NAND do it? Samsung is keeping quiet. It's known to be a derivative of SLC NAND, perhaps with shorter bitlines and wordlines. With so little to go on, it's hard to be sure how the competition will play out. It's not even clear the extent to which Z-NAND is a proprietary tech based on intellectual property that can't legally be copied.

That's the greatest worry regarding 3D Xpoint—that Intel will slow-ball development if it ever gains a significant advantage. But even if Z-NAND does end up being a Samsung-only tech, Samsung versus Intel will drive the industry forward faster than just one of them dominating the market for next-gen non-volatile memory, that's for sure.

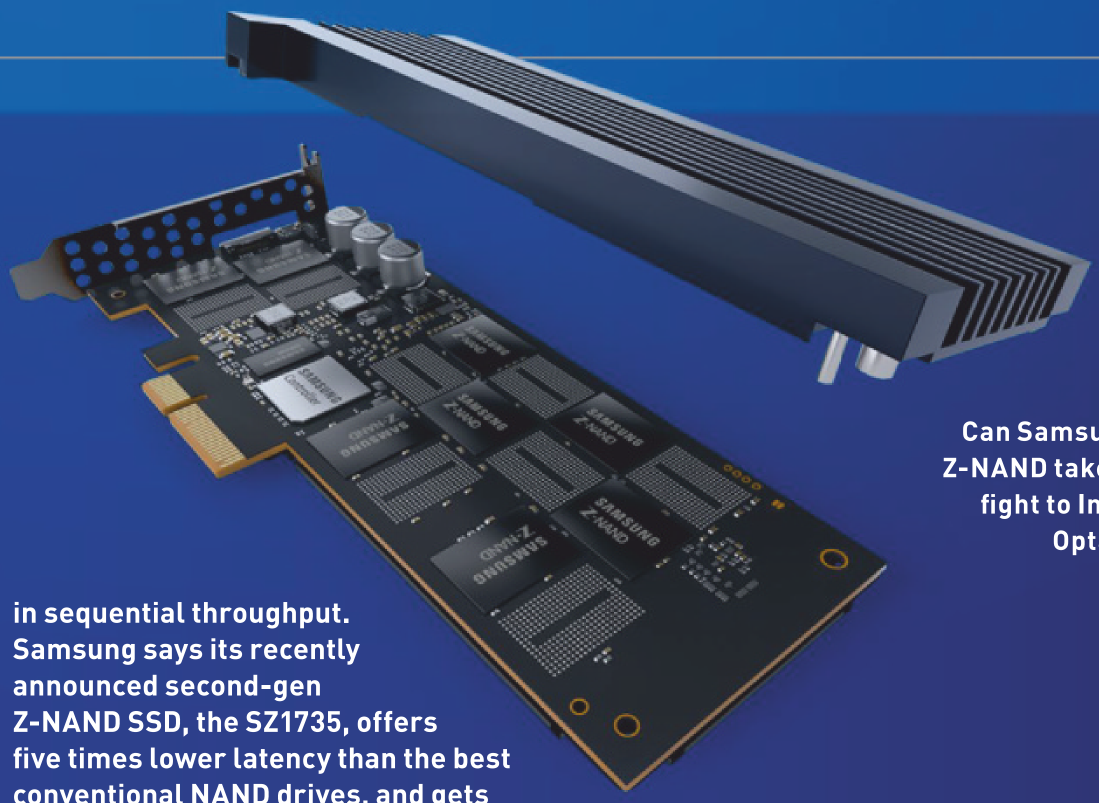

Can Samsung's Z-NAND take the fight to Intel's Optane?

The home of technology

techradar.com

SUBSCRIBE TO MAXIMUMPC

THE GO-TO RESOURCE FOR PC ENTHUSIASTS...

CHOOSE YOUR PERFECT PACKAGE

AMD RADEON RX 590 The affordable GPU for today's gaming PG. 74
SOLID-STATE NATION Find the perfect SSD for your system & budget PG. 36
ACER PREDATOR X27 Is this the best screen you can buy right now? PG. 82
MAXIMUMPC
2019 TECH PREVIEW
Next-gen hardware from AMD, Intel, & Nvidia that will transform your PC
THE BEST GAMES COMING THIS YEAR Revealed: The biggest & finest titles for 2019 PG. 48

THREADRIPPER 2970WX AMD's new workstation CPU rips up the benchmarks PG. 74
GEFORCE RTX 2070 FE Nvidia's mid-range ray-tracer hits the testing labs PG. 76
DRIVE IMAGING 101 Upgrade to a speedy SSD the fast and easy way PG. 48
MAXIMUMPC
2019 TECH PREVIEW
Next-gen hardware from AMD, Intel, & Nvidia that will transform your PC
PLUS! Build a 1080p gaming rig PG. 68
THE BEST GAMES COMING THIS YEAR Revealed: The biggest & finest titles for 2019 PG. 48

1

THE DIGITAL EDITION

ONLY $9.00

EVERY YEAR

Instant digital access on your iPad, iPhone, and Android device.

2

THE PRINT EDITION

ONLY $24.00

EVERY YEAR

Every issue delivered to your door for a fraction of the store price.

3

COMPLETE PRINT & DIGITAL BUNDLE

ONLY $28.00

EVERY YEAR

Get the print edition of *Maximum PC* delivered direct to your door, and instant access on your iPad, iPhone, or Android device.

MAXIMUM PC
GEAR of the YEAR
2019 TECH PREVIEW
Next-gen hardware from AMD, Intel, & Nvidia that will transform your PC
THE BEST GAMES COMING THIS YEAR
PLUS! Build a 1080p gaming rig PG. 68

On iOS & Android!

THE EASY WAYS TO SUBSCRIBE…

http://bit.ly/sub-maxpc

Or **call 844-779-2822** (toll-free)

TERMS AND CONDITIONS Prices and savings quoted are compared to buying full-priced US print and digital issues. You will receive 13 issues in a year. You can write to us or call us to cancel your subscription within 14 days of purchase. Your subscription is for the minimum term specified, and will expire at the end of the current term. Payment is non-refundable after the 14-day cancelation period, unless exceptional circumstances apply. Your statutory rights are not affected. Prices correct at time of print and subject to change. Call center opening hours are Monday to Friday 8am–7pm, and Saturday 10am–2pm EDT. For full terms and conditions, please visit http://bit.ly/magtandc. **Offer ends: February 28, 2019.**

TECH PORN

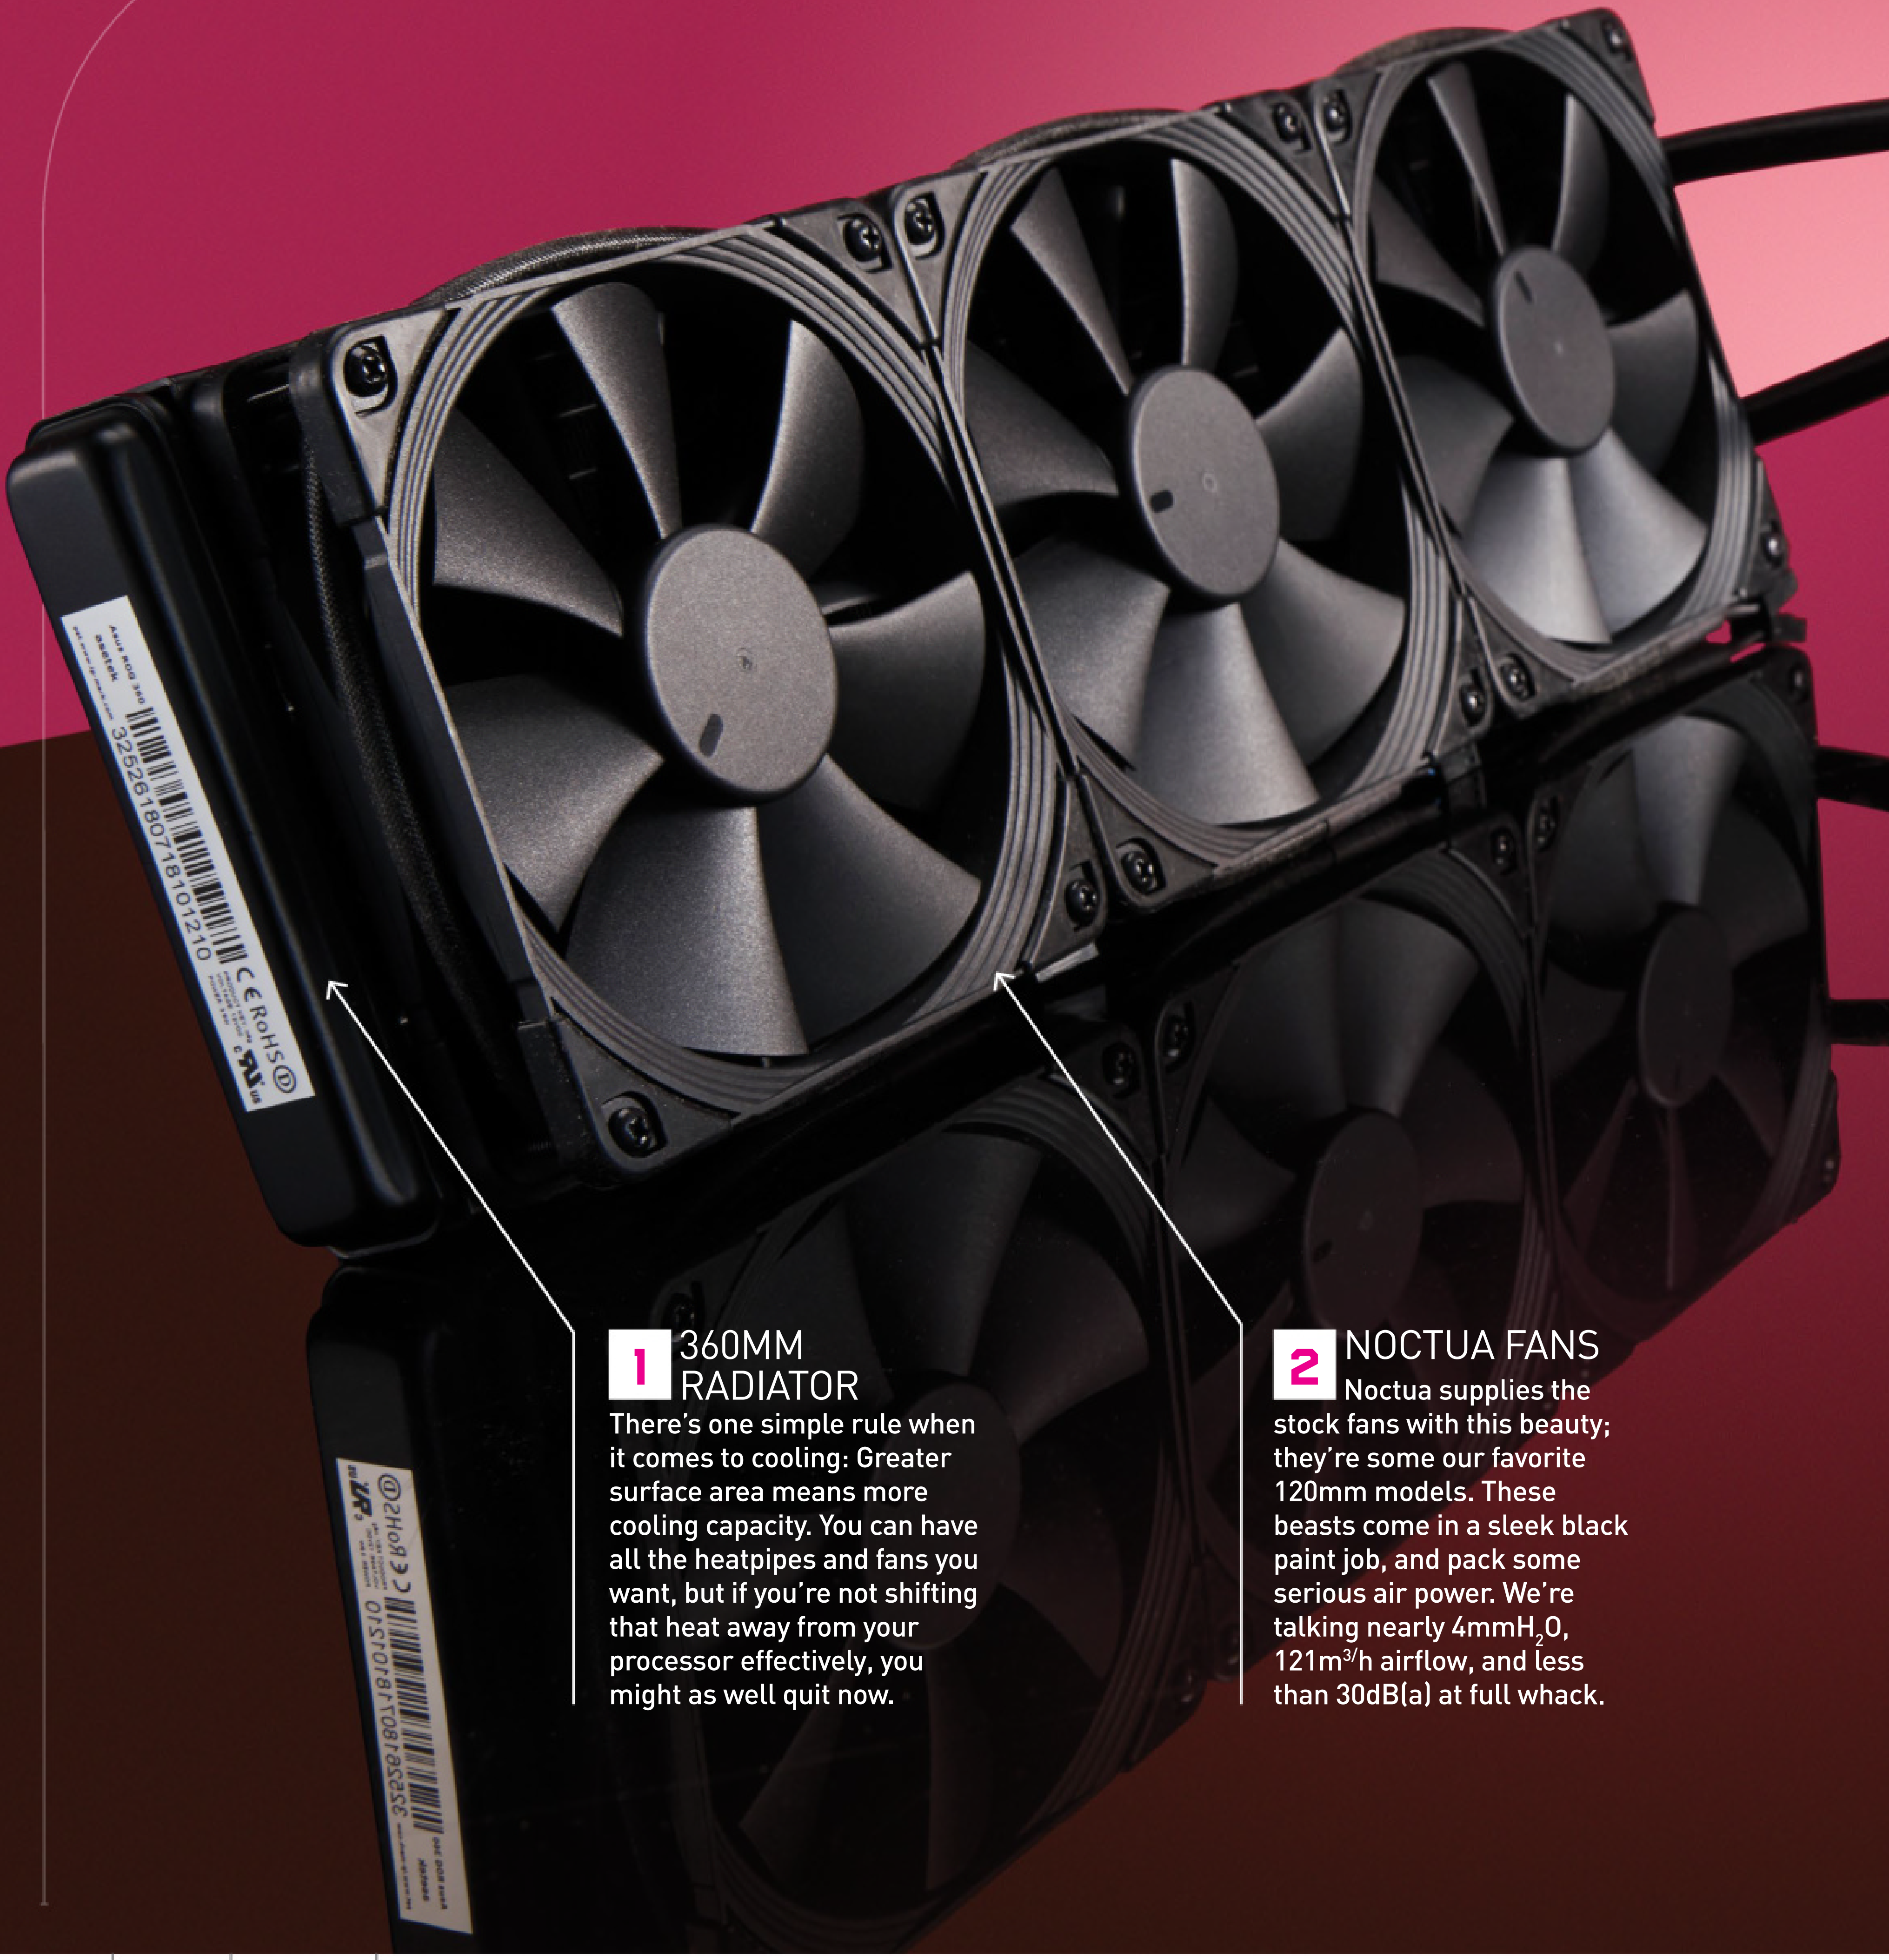

1 360MM RADIATOR

There's one simple rule when it comes to cooling: Greater surface area means more cooling capacity. You can have all the heatpipes and fans you want, but if you're not shifting that heat away from your processor effectively, you might as well quit now.

2 NOCTUA FANS

Noctua supplies the stock fans with this beauty; they're some our favorite 120mm models. These beasts come in a sleek black paint job, and pack some serious air power. We're talking nearly 4mmH_2O, 121m^3/h airflow, and less than 30dB(a) at full whack.

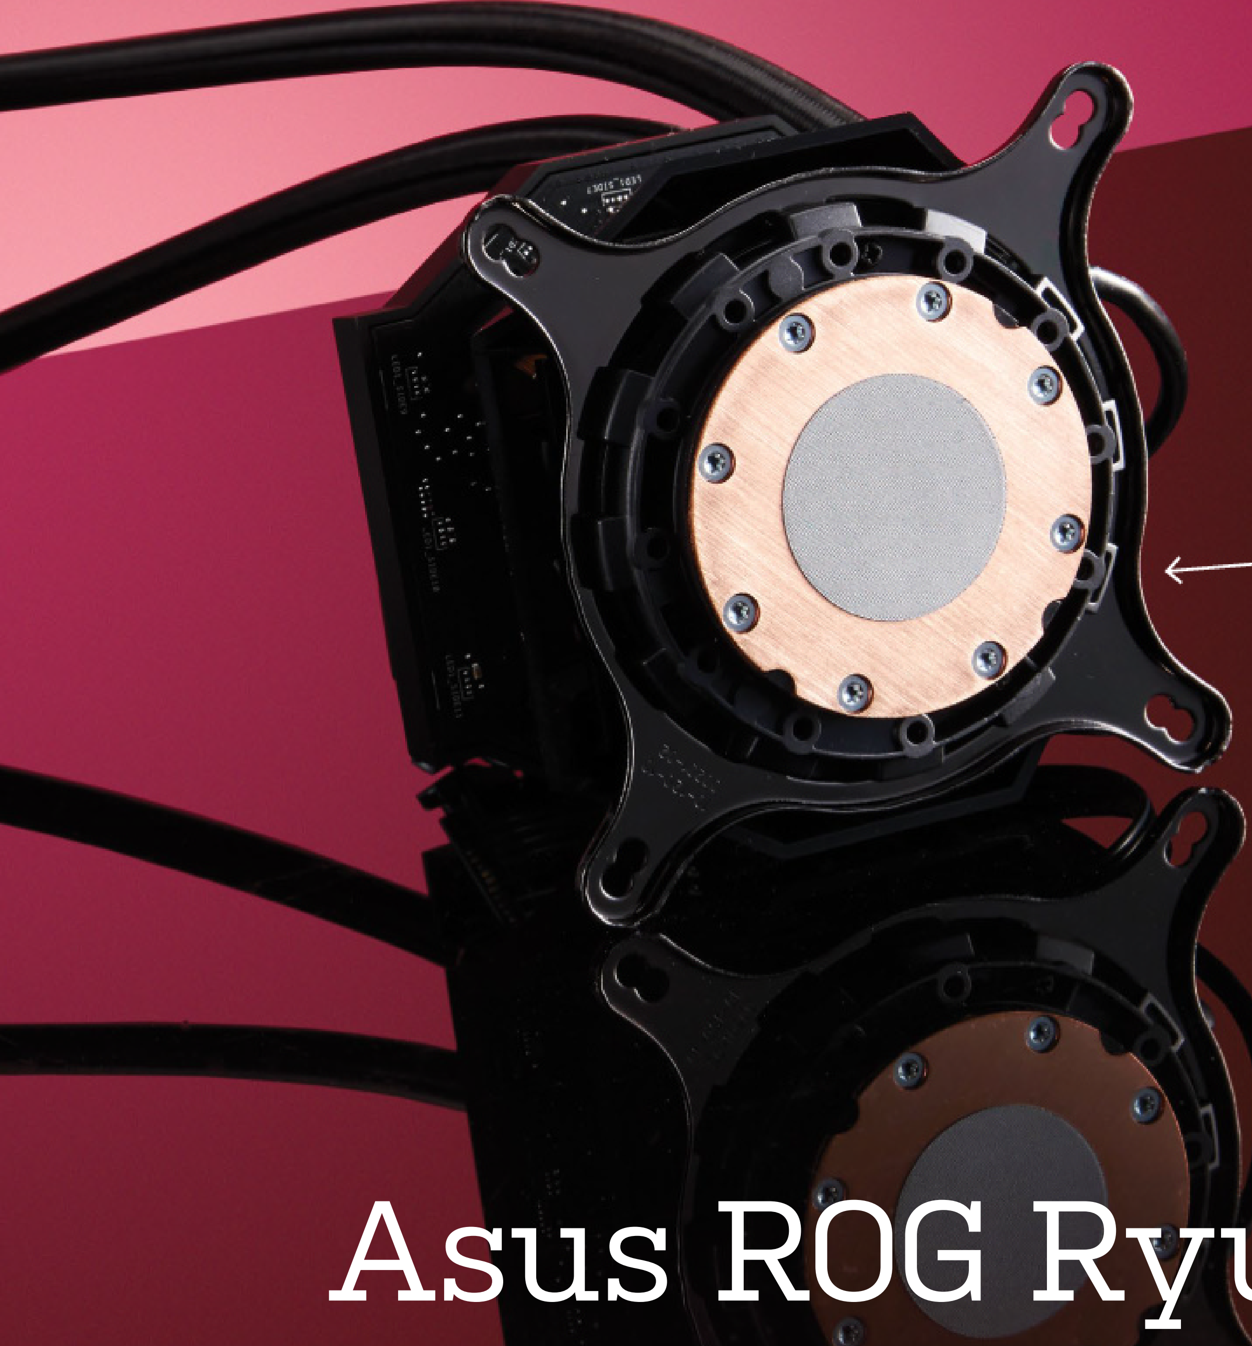

3 OLED DISPLAY FOR DAYS

The showpiece (and clearly an Asus trend lately) is the inclusion of the 1.77-inch OLED display. This sits on top of the Asetek round pump we all know and love, and is fully addressable, meaning you can stick any number of animations, images, or system stats on it.

Asus ROG Ryujin 360

LET'S FORGET ABOUT the name for a second—we know it's ridiculous. Seriously, who gives a product a name like that? We keep having to refer back to the documentation to check how to spell it: R-yu-jin. But let's take a moment to consider what Asus has pulled off here. Cooling: Yeah, it might be an Asetek-designed pump and mounting mechanism, and, sure, there isn't any of that fancy Threadripper-sized CPU block to cover those gargantuan dies, but the innovations included in this—where do we even begin?

From the Noctua industrial fans to the OLED display baked into the CPU block cover to the black sleeved tubing, it's all there. For a first attempt at an AIO, Asus really is batting it out of the park.

The world of all-in-one cooling tech might well have originated from the likes of Corsair and Cooler Master, but it's currently companies such as NZXT and Asus who are really pushing the boundaries when it comes to innovation in this field. Do you need a 360mm triple radiator for your rig? Probably not. Do you need a fancy OLED display for your system? Hmm, no. Does it look cool? Darn right it does. Still think that name is bonkers, though....

GAMES FOR 2019

From heavy gunfire to forest spirits, here are our top tips for 2019

BY IAN EVENDEN

SEQUELS AND REMAKES, SEQUELS AND REMAKES. It's been a criticism of the games industry for some time that it's averse to taking a risk on something new, more content to rake over the past, and present the familiar and "safe." This has been particularly true ever since the idea of the "HD remaster" came into being, with old games being polished up for modern graphical standards—to varying degrees of success—before being set before consumers as something new.

There is innovation—if you know where to look. The coming year sees a raft of Japanese games arrive on PC, the sort of thing that once would have stayed on PlayStation, while the unstoppable march of *Warhammer 40K* continues, and Tom Clancy is still producing games despite being dead for five years.

The PC is in a strong place as a games machine—the consoles look exhausted as they gear up for another generational leap—and it's the place to come to find the most spectacular graphics and the cleverest use of AI, as this brief rundown of games to look out for in 2019 shows. There are a few sequels and remakes, too, though.

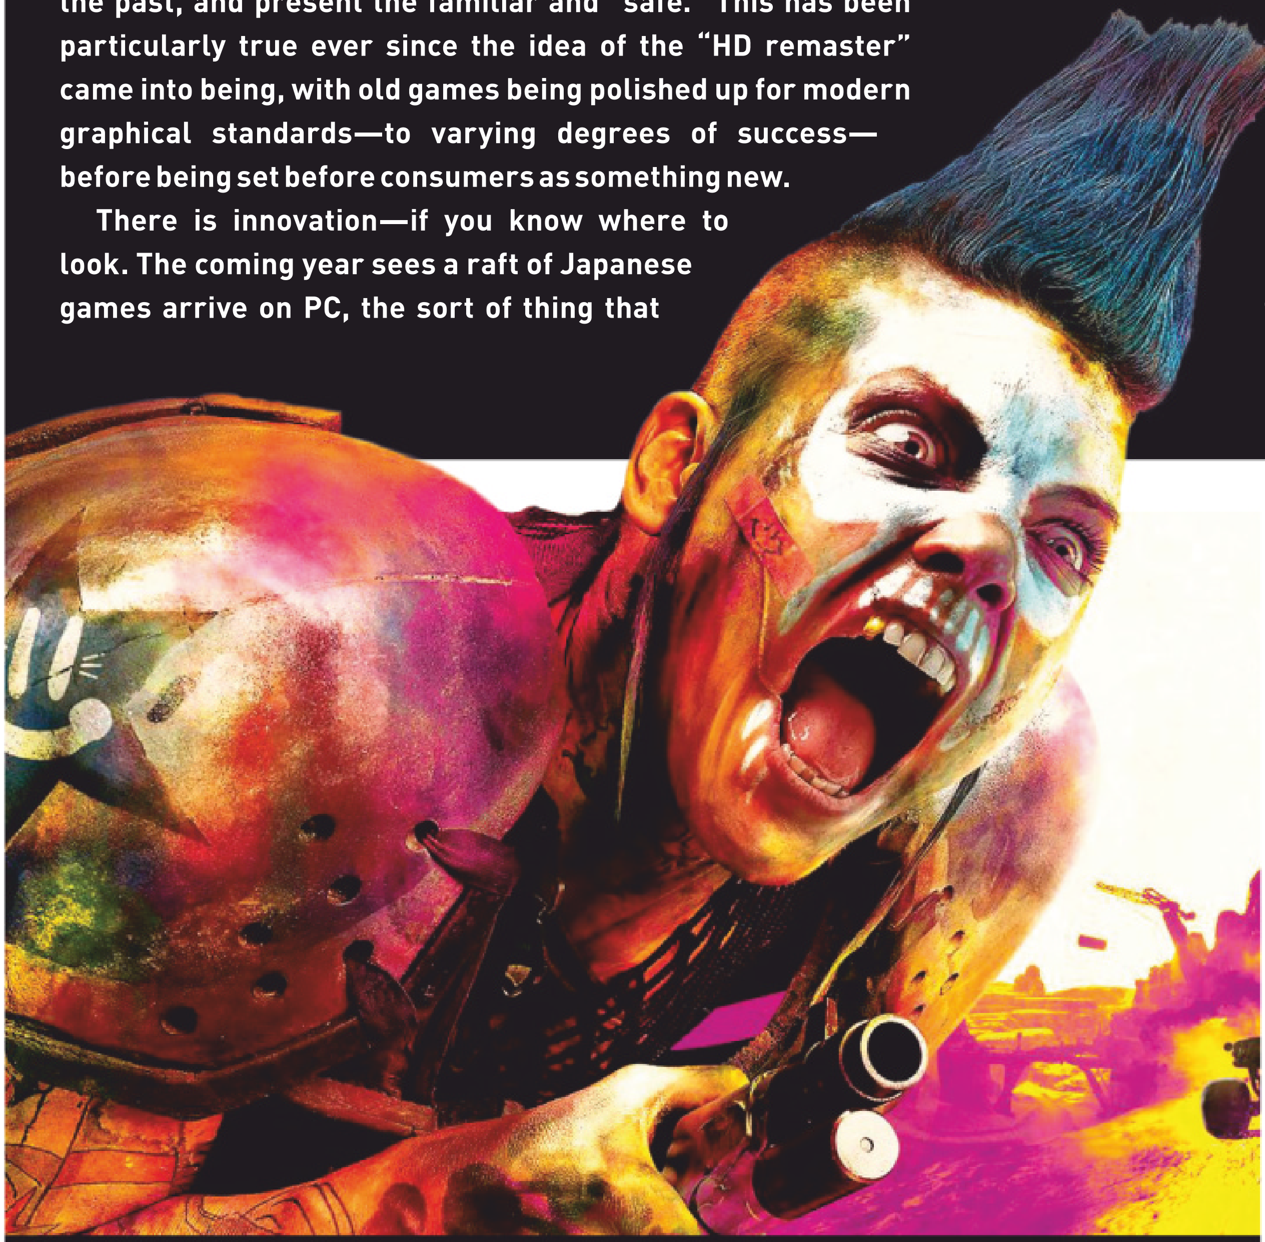

Rage 2 TBC 2019

The sequel to id Software's post-apocalyptic FPS moves to Avalanche, home of *Just Cause,* where it picks up an open-world environment and a heavy nod to *Fury Road.* Once again, you're shooting mutants and driving messed-up vehicles, as well as lopping off heads with a razor-sharp boomerang, but this time it's using actual colors, rather than gray and brown.
» **Publisher:** Bethesda » **Developer:** Avalanche Studios

Dirt Rally 2 02/26/19

The official game of the FIA World Rallycross Championship, *Dirt Rally 2* will take you to six rally locations around the world. The *Dirt* games have always been fun—there's a primal thrill in hurling a howling machine down muddy tracks and up gravelly mountains—and there's no reason to think this will be any different.
» **Publisher:** Codemasters
» **Developer:** Codemasters

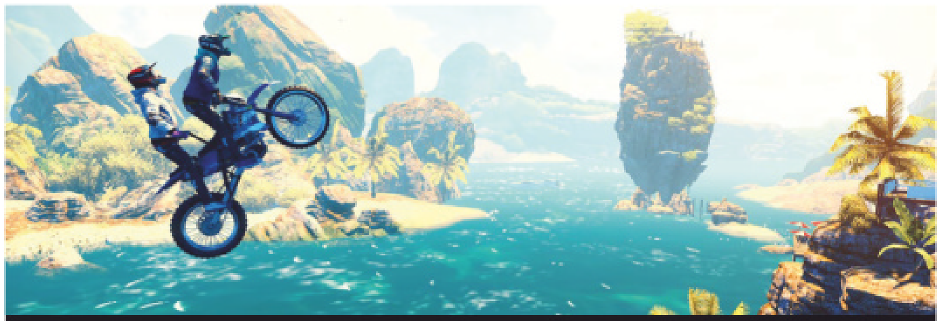

Trials Rising

02/12/19

This ragdoll physics bike racer is poised to smash its way through unconvincing scenery following its beta at the end of 2018, bringing a new world-tour backdrop that plots courses around landmarks such as the Eiffel Tower. A new co-op mode will see two players control a single bike—a recipe for disaster if we ever heard one.

» **Publisher:** Ubisoft » **Developer:** RedLynx/Ubisoft Kiev

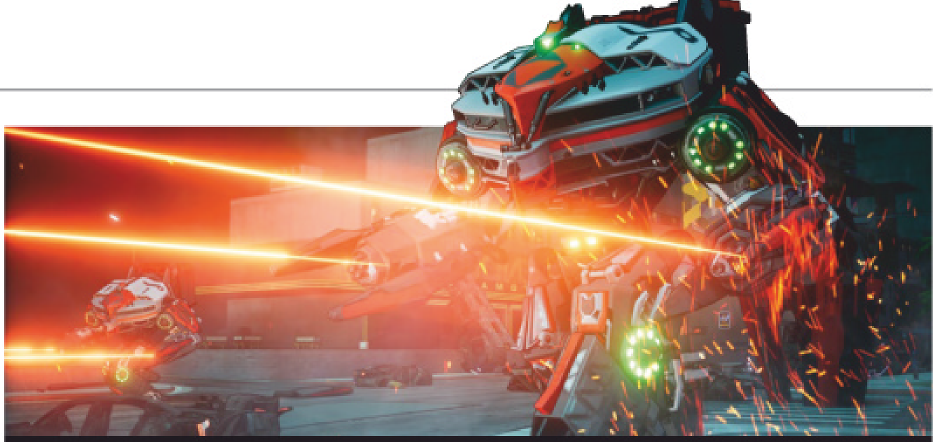

Crackdown 3

02/15/19

The Agency, a group of superpowered madmen, who've been known to drive jeeps straight up vertical buildings and blow the heads off hundreds of goons, is called into action after a terror attack takes out the power all around the world. We reckon we know what happens next, and it's not going to involve wiring puzzles.

» **Publisher:** Microsoft » **Developer:** Sumo Digital

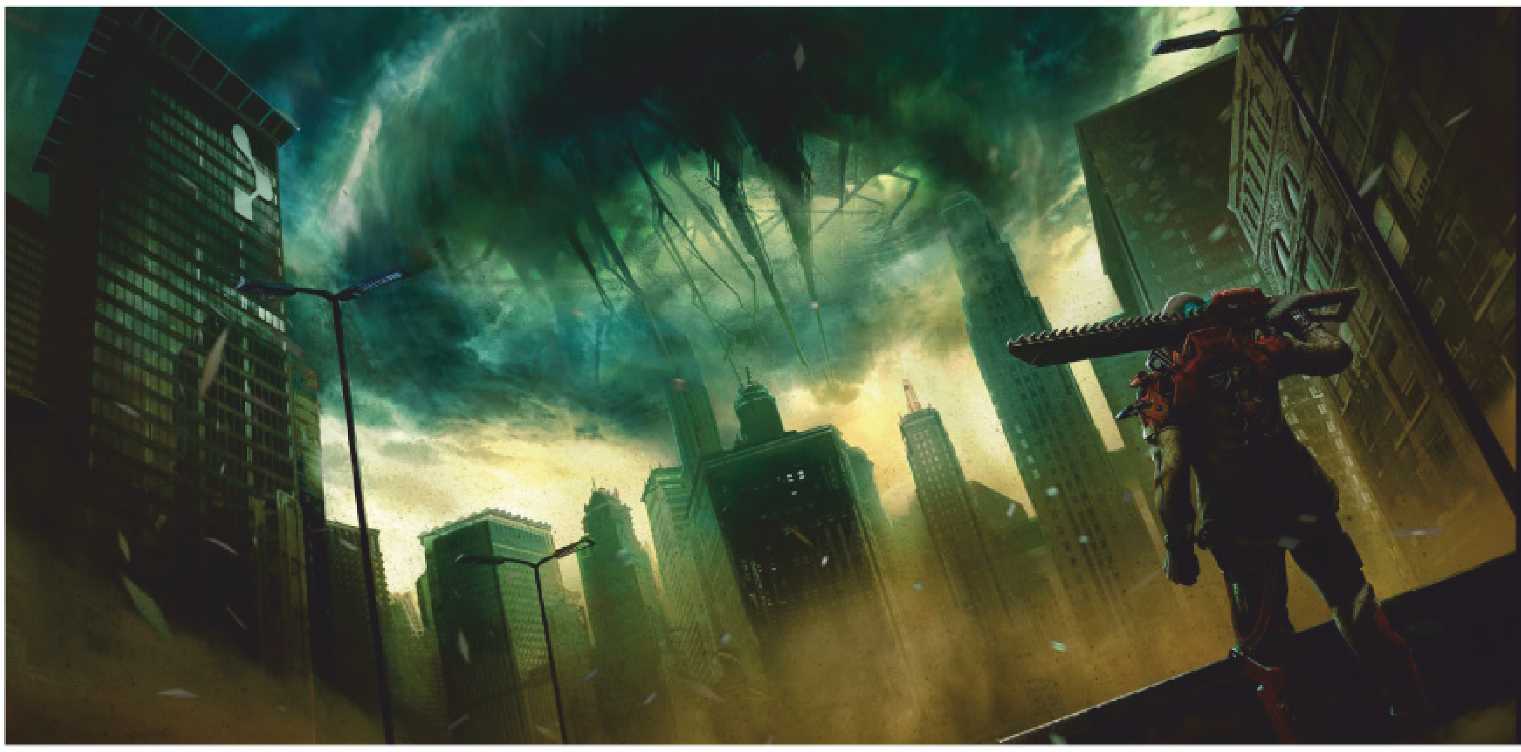

The Surge 2

TBC 2019

A sci-fi take on the *Souls* games was the USP of the original release, and while *The Surge* didn't quite get the acclaim that FromSoftware's series has, it was definitely worth picking up for a sequel. An open-world action RPG in which you fight robots—you can expect oil to be splattered everywhere.

» **Publisher:** Focus Home Interactive

» **Developer:** Deck 13

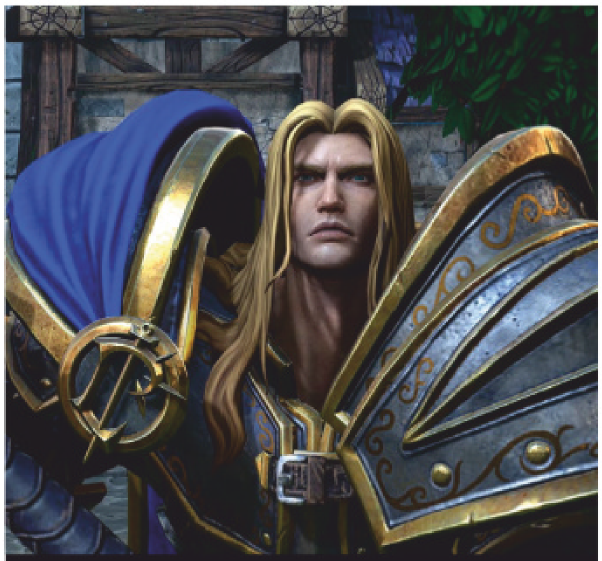

Warcraft III: Reforged

TBC 2019

A remake of the 2002 RTS original made with the events of *WOW* in mind, it looks stunning. Hopefully, the retconned events and rebalanced gameplay will be as much of a time sink as first time. The original will get a patch, too, and owners will be able to play against *Reforged* players.

» **Publisher:** Blizzard Entertainment

» **Developer:** Blizzard Entertainment

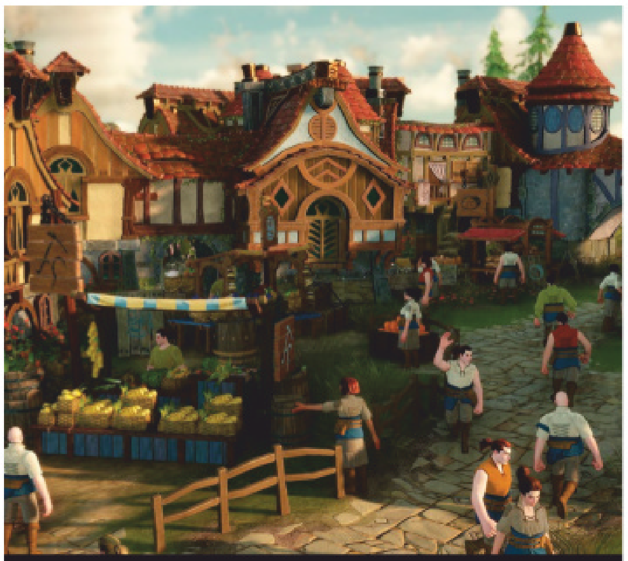

The Settlers

Fall 2019

German dev Blue Byte has picked up the reins of its venerable RTS series and completely updated it for modern PCs and Ubisoft's Snowdrop engine, so it's looking quite splendid. And if you're hot for some medieval town-building, there's an updated version of the original game on uPlay, too.

» **Publisher:** Ubisoft

» **Developer:** Blue Byte

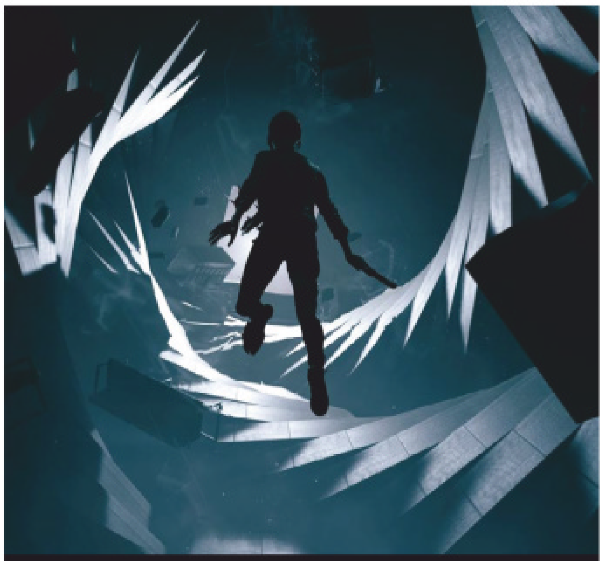

Control

TBC 2019

Alan Wake dev Remedy is behind this third-person adventure in which a psychic, perhaps supernatural, woman fights her way through a brutalist skyscraper with a shape-changing pistol. After *Alan Wake's American Nightmare,* we suppose this was inevitable.

» **Publisher:** 505 Games

» **Developer:** Remedy Entertainment

Sekiro: Shadows Die Twice 03/22/19

Coming from FromSoftware, the studio behind *Dark Souls*, *Armored Core*, and *Bloodborne*, we're not expecting this to be an easy game. Set in Sengoku Japan, and following a mysterious one-armed warrior, *Sekiro* eschews the usual health bars in favor of one-hit kills, but you must knock your opponent off balance first in a graceful duel.

» **Publisher:** Pub Activision
» **Developer:** FromSoftware

Devil May Cry 5 03/08/19

Not a traditional PC franchise, but the boundaries are dissolving now there's X86 at the heart of consoles, too. This one's about demon hunters, one of whom uses his motorbike and hat as weapons. There are also a lot of guns and swords, occasionally a massive fleshy scythe, and their arms come off. Very confusing, but a lot of fun, and spectacular in motion.

» **Publisher:** Capcom
» **Developer:** Capcom

Anno 1800 February 2019

The *Anno* universe hits the industrial revolution, and sees you use the power of metal and steam engines to build a rural community into an industrialized city. The series has melded elements of grand strategy, such as diplomacy and warfare, on to its city-building core, so seeing what it can do with this tumultuous time in history will be extremely interesting.

» **Publisher:** Ubisoft
» **Developer:** Blue Byte

Phoenix Point June 2019

From one of the men behind the original *X-COM* (that's 1994, not 2012) this is a spiritual successor to the series—so we know how it's going to play out. Turn-based combat, an alien invasion, and a resistance cell bravely battling both monsters and resource management trying to take the bad guys down. Frankly, we can't wait.

» **Publisher:** TBC
» **Developer:** Snapshot Games

Age of Wonders: Planetfall TBC 2019

Turn-based empire building in a sci-fi setting with multiple paths to victory sounds great, and when it comes from the dev behind *Age of Wonders* dating back to 1999, we like it even more. Bought by Paradox last year, the Dutch studio should now have the backing and resources it needs to knock this into orbit.

» **Publisher:** Paradox Interactive
» **Developer:** Triumph Studios

Resident Evil 2 01/25/19

A remake of the 1998 PlayStation classic, this has learned from 20 years of game releases, and implemented things such as a checkpoint save system, controls that don't make you want to scream, and a proper PC release. Otherwise, it's exactly what you'd expect, as an army of the undead is fought off by two characters using guns, grenade launchers, and lockpicks.

» **Publisher:** Capcom » **Developer:** Capcom

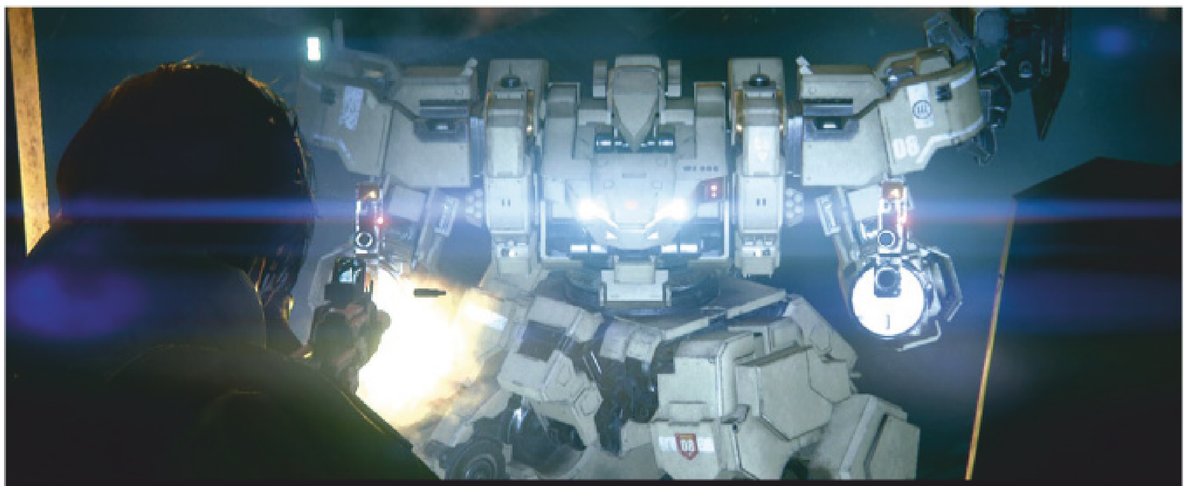

Left Alive 03/05/19

A shooter set in the *Front Mission* universe—that means enormous stompy robots—*Left Alive* spins a tale of honor, revenge, and despair. And huge metal men. We'd be disappointed if you didn't get to blow at least one up, but the trailers so far have been heavy on characterization and *Metal Gear*-like tactical shooting.

» **Publisher:** Square Enix » **Developer:** Square Enix

Skull & Bones

TBC 2019

Open-world tactical action tends to bring up images of gruff men in black with guns—but this one is about pirates. The characters may still be gruff, but there's a chance of a shanty as you either pursue the single-player campaign or team up with five others to battle against other ships.

» **Publisher:** Ubisoft

» **Developer:** Ubisoft Singapore

Tom Clancy's The Division 2 03/15/19

Another persistent-world multiplayer from Ubisoft, this time in Washington DC in summer. The viral apocalypse still took place, but you have new weapons and skills to deal with the enemies it caused.

» **Publisher:** Ubisoft

» **Developer:** Massive Entertainment

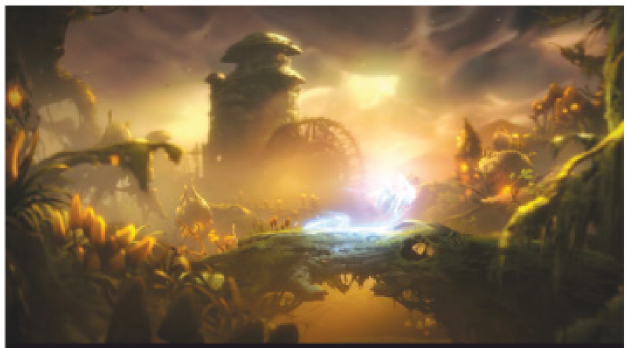

Ori and the Will of the Wisps TBC 2019

Ori and the Blind Forest was a charming hit in 2015, and this promises more of the same. A guardian spirit moves between platforms and solves puzzles, while getting stronger and unlocking new areas.

» **Publisher:** Microsoft

» **Developer:** Moon Studios

Life is Strange II

Episode 1 out now, others TBC

Taking its structure from road movies, Sean is on the run from the police with his brother—who has mysterious powers—after their father was killed in what appears to be a supernatural incident.

» **Publisher:** Square Enix

» **Developer:** Dontnod Entertainment

Shenmue III 08/27/19

Really? Yes, it seems it actually will come out. This "for the fans, by the fans" is actually by series creator Yu Suzuki, and sees martial artist Ryo journey into the Japanese mountains, where he meets a mysterious girl and a new enemy.

» **Publisher:** Deep Silver
» **Developer:** Ys Net/Neilo

Psychonauts 2 04/19/19

Here's something we don't see much of: a platform game. This crowdfunded sequel sees Raz and his psychic spies explore the mental landscapes of in-game characters, which become the game's levels, to be traversed using a range of powers.

» **Publisher:** Majesco Entertainment
» **Developer:** Double Fine

Tropico 6 January 2019

The bonkers Caribbean dictator sim is back for another grasp at power, bringing large archipelagos to the series. You also send agents to other countries to steal monuments, make promises you can't keep, and gold-plate your palace.

» **Publisher:** Gathering of Developers
» **Developer:** PopTop Software

Gears 5

TBC 2019

Though no longer officially "Of War," the Gears are still fighting, now in a quest to discover the origin of the Locust Horde and, presumably, carve it in half with a chainsaw. JD Fenix, the protagonist of gaming's least subtle third-person blaster, returns, as does his dad Marcus, who's still voiced by Bender.

» **Publisher:** Microsoft
» **Developer:** The Coalition

Wolfenstein: Youngblood TBC 2019

MachineGames clearly wanted a change; this co-op game is set in the 1980s, and stars BJ's twin girls. The Nazis are still in power in some places, and need to be taken out with extreme prejudice. Goody!

» **Publisher:** Bethesda
» **Developer:** MachineGames

A Plague Tale: Innocence TBC 2019

Children struggling to survive the Hundred Years War is a jolly premise. Plus, they're stalked by hungry rats, which they must drive off with fire or lure into the path of knights. One to watch. From a distance.

» **Publisher:** Focus Home Interactive
» **Developer:** Asobo Studio

Battlefleet Gothic: Armada 2 01/24/19

Smash huge spaceships into even bigger biological terrors in this *Warhammer 40,000* strategy. Twelve factions will be playable at launch, and ships eating other ships aren't out of the question.

» **Publisher:** Focus Home Interactive
» **Developer:** Tindalos Interactive

Ace Combat 7: Skies Unknown 02/01/19

In the absence of *Top Gun 2*, this will have to do. *Ace Combat* is a series about hurling millions of dollars of jet around the skies and not caring too much. It's definitely not a simulation, because we're fairly sure the USAF doesn't do this many loops in normal operations, but it is fast and spectacular, with lots of gratifying explosions.

» **Publisher:** Bandai Namco
» **Developer:** Bandai Namco

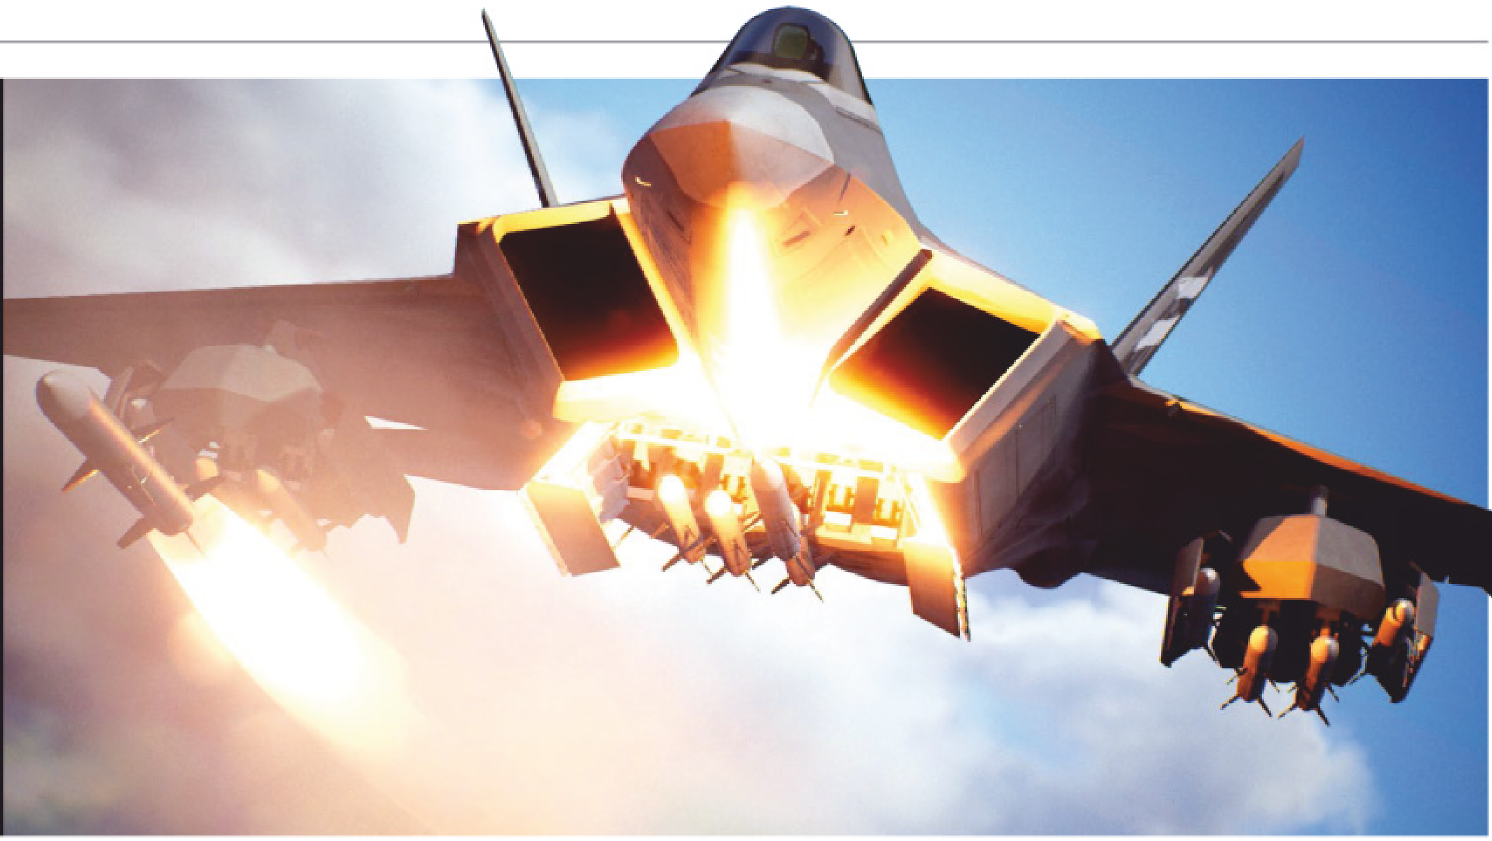

Metro Exodus 02/22/19

You know, like when everyone tries to get off the subway at once, only with even more guns. The third *Metro* game starts off in a harsh, post-apocalyptic winter, before sending protagonist Artyom out of the city and across an entire continent. No doubt this break from the claustrophobic norm will be interrupted by bitter fighting and mutated creatures, as usual.

» **Publisher:** Deep Silver » **Developer:** 4A Games

Total War: Three Kingdoms 03/07/19

Creative Assembly takes its historical battle series to China, where 11 factions are fighting it out to unify the country into the behemoth we know today. Romance mode, based on the classic novel *Romance of the Three Kingdoms*, will gift generals with superhuman powers, while in other modes those generals will form social connections and relationships. Before spilling lots of blood, of course.

» **Publisher:** Sega » **Developer:** Creative Assembly

SPQR
12 Apr, 450
Legion I
8,000
32,694
Corfinium's 1ST Cohort
Fregellae's 1ST Cohort
Venafrum's 1ST Cohort
Teanum Apulum's 1ST Cohort
Nola's 1ST Cohort
Nola's 2ND Cohort

Imperator: Rome TBC 2019

More strategy from Paradox, this time with a sequel to 2008's *Europa Universalis: Rome*, and covering the establishment of the Roman Empire 2,500 years ago. While most of the game will see you battling on a grand scale, there will also be character management, in the style of *Crusader Kings II*.

» **Publisher:** Paradox Interactive
» **Developer:** Paradox Development Studio

gamesradar+

One-stop destination for games, movies, TV, and more!

New Channels: Movies, TV and Cool Stuff!

www.gamesradar.com

HOW TO

STEP-BY-STEP GUIDES TO IMPROVING YOUR PC

TIP OF THE MONTH

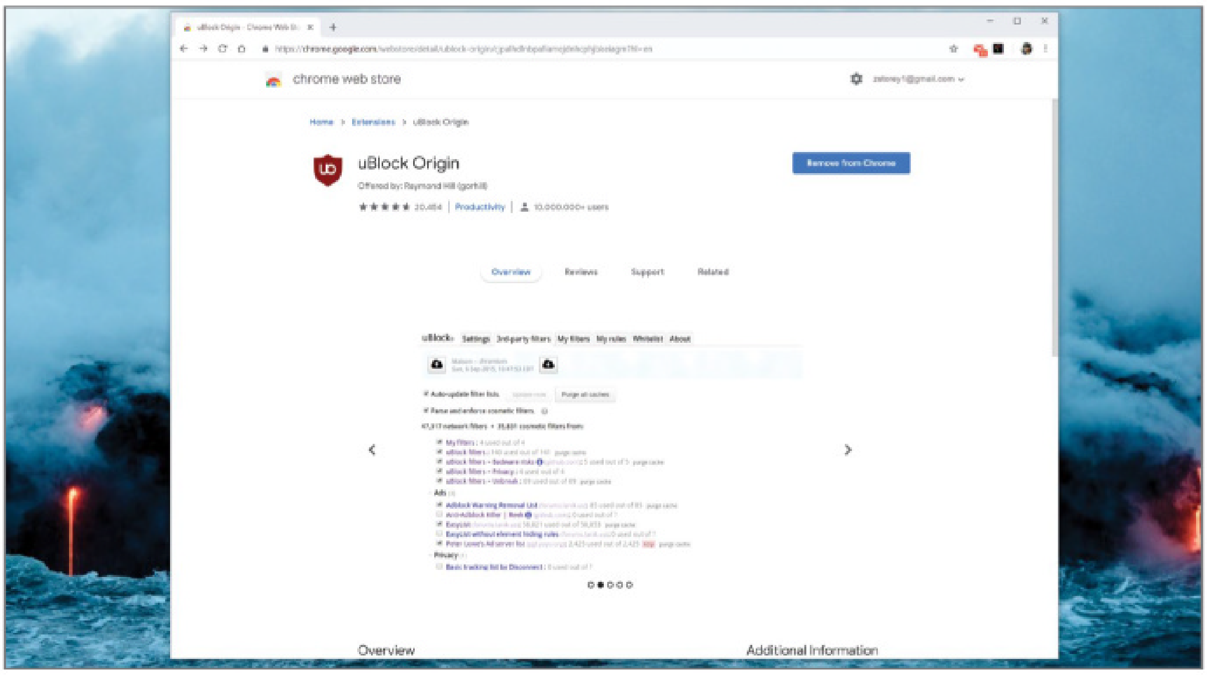
chrome web store
Home > Extensions > uBlock Origin
uBlock Origin
Offered by: Raymond Hill (gorhill)
20,484 | Productivity | 10,000,000+ users
Remove from Chrome
Overview
Reviews
Support
Related
uBlock
Settings
3rd-party filters
My filters
My rules
Whitelist
About
Overview
Additional Information

THE BEST AD-BLOCKER: UBLOCK ORIGIN

Shh, don't tell anyone, but we don't like those pesky ads either. Let's face it, relevant ads are good, but ones that take over an entire page or drag you elsewhere aren't good. uBlock Origin is an open-source, low-resource-impacting ad-blocker, which is frequently updated—and we're big fans. Add filters for your favorite content producers, and enjoy super-snappy load times on everything else. http://bit.ly/UBlockOrigin $free

MAKE – USE – CREATE

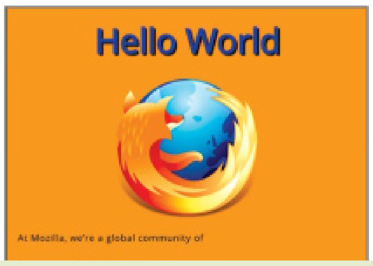
Hello World
At Mozilla, we're a global community of

58
Program and set up a tiny web server with Linux

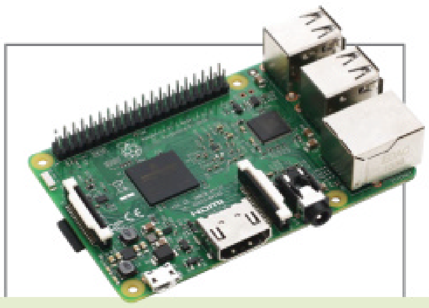

62
Build a Raspberry Pi-powered random video player

68
How we built this month's Full HD AMD gaming rig

ZAK STOREY
DEPUTY EDITOR

OH, MSI, WHY?

There are two types of people: the organized, who have everything clean, spick and span, and in the right place; and the chaos curators, who have everything in a mess, but kinda know where it's all located. The latter, at least for now, is my life. As I'm finally moving into my first home, the word "pandemonium" comes to mind.

Time is short, and I find myself in bed each night wondering how I got there, and where the day's gone. So when MSI's Command Center threw up a notification on my rig telling me I had updates to install, I thought, why not?

What a mistake that was. I clicked "Yes," MSI showcased a list of legit driver updates, I clicked "Continue," and MSI finished up. I came back a moment later, and BAM! McAfee popped up: "Your antivirus has expired and you haven't run a scan ever."

That's interesting, as I've never used McAfee on this system. In fact, I avoid all antivirus outside of Windows Defender. So, how the hell did McAfee get on my rig? The only conclusion I can make: MSI's Command Center installed it on the sly.

Sure enough, I opened "Programs and Features," and sitting there in the application list was McAfee, installed that day. I promptly removed it, and got the usual garbage. "Why am I leaving McAfee?" the customer feedback survey asks. "Because I don't want bloatware, stealthily installed on my rig, and if I did need antivirus, I'd go with Avira or Bitdefender."

↘ submit your How To project idea to: comments@maximumpc.com

IFIXIT presents:

AUTOPSY

THIS MONTH WE DISSECT...

Mac Mini (2018)

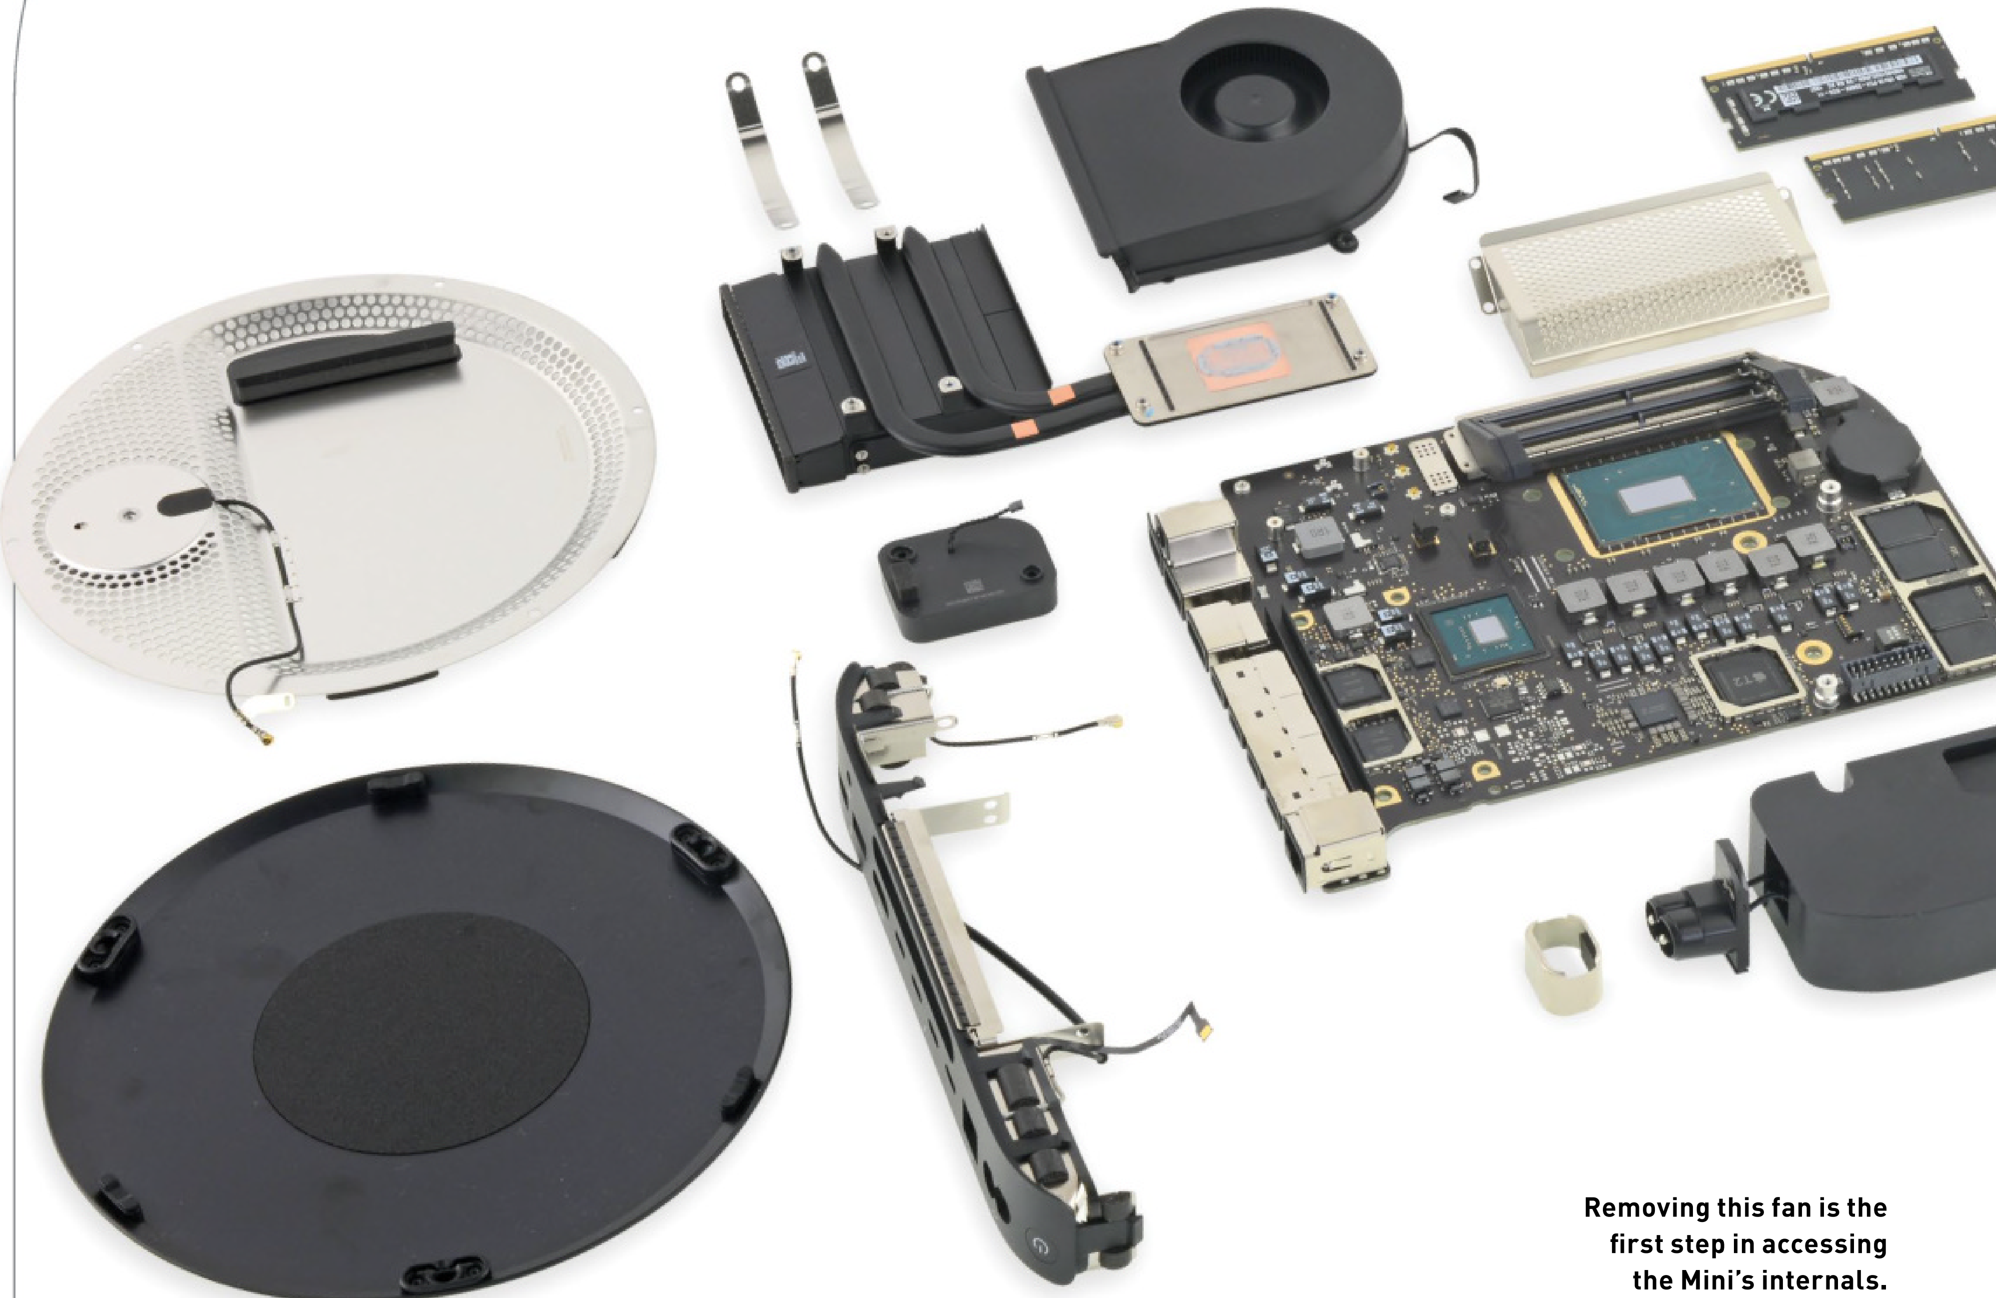

Removing this fan is the first step in accessing the Mini's internals.

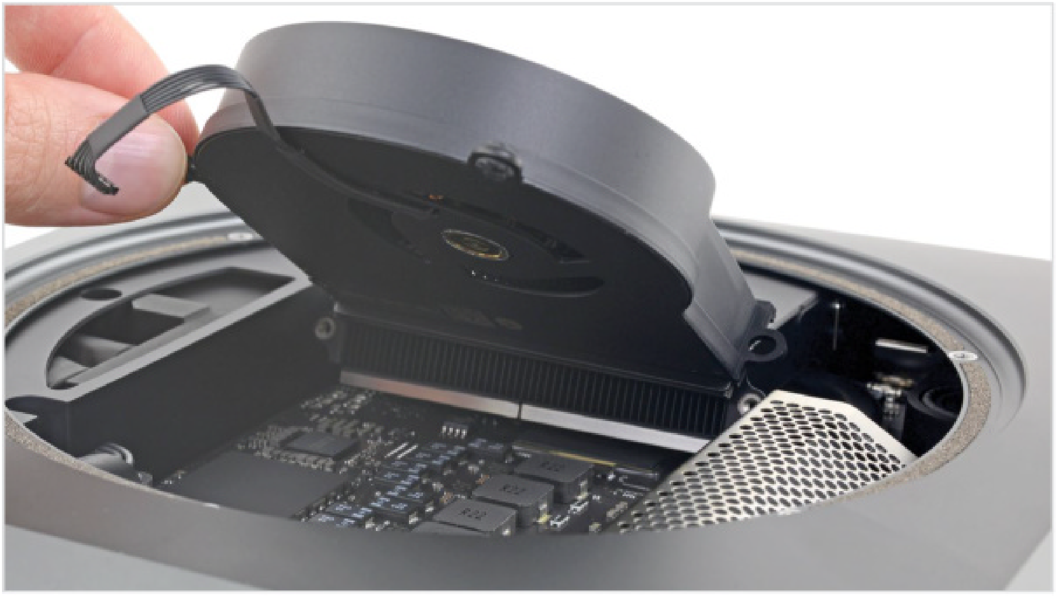

About iFixit

iFixit is a global community of tinkerers dedicated to helping people fix things through free online repair manuals and teardowns. iFixit believes that everyone has the right to maintain and repair their own products. To learn more, visit www.ifixit.com.

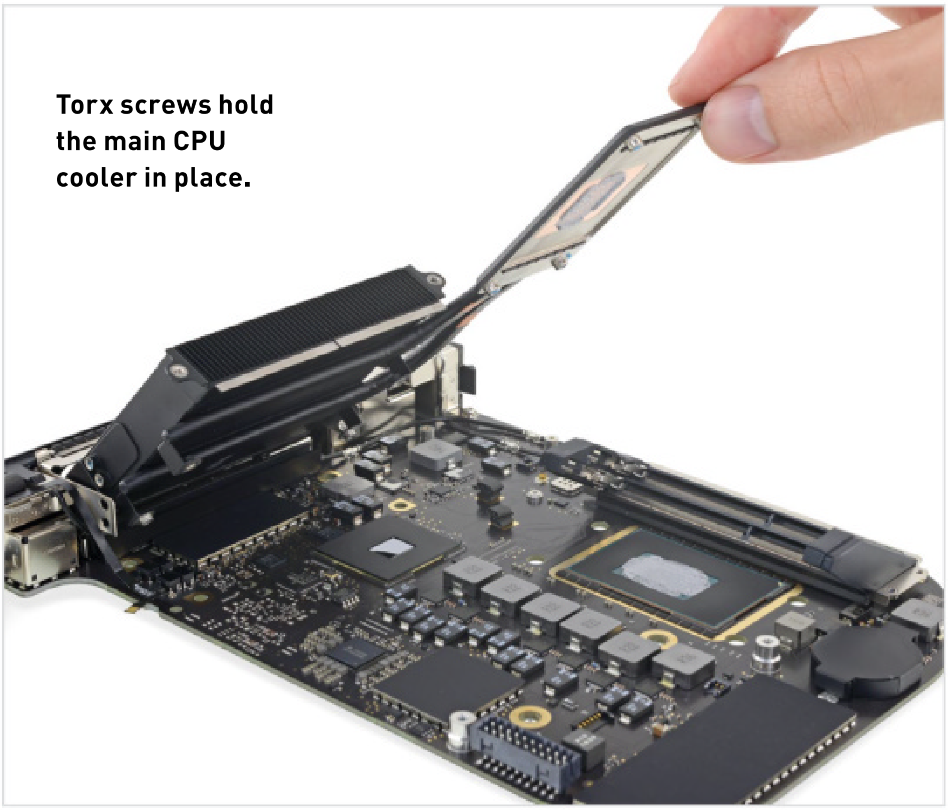

Torx screws hold the main CPU cooler in place.

BACKGROUND

Apple's once-neglected Mac mini is coming in hot with a brand-new, cutting-edge, long-awaited... processor upgrade? And a couple more ports? There must be more, and we know how to find it—with a teardown!

MAJOR TECH SPECS

- 3.6GHz quad-core Intel Core i3 with 6MB shared L3 cache
- 8GB of 2,666MT/s DDR4 SO-DIMM memory
- 128GB SSD
- Intel UHD Graphics 630
- 802.11ac Wi-Fi and Bluetooth 5.0
- macOS Mojave

KEY FINDINGS

- Our first look at the 2018 mini's exterior gives us the warm and fuzzies—it's the same friendly form factor we remember. Despite controversial departures from a few common ports, Apple has included plenty of them here. We spot two USB-A ports, four USB-C ports, a headphone jack, an Ethernet port, and an HDMI port (which isn't available on any other recent Apple product).

- With some trepidation, we point our tools at the 60-percent-recycled-plastic bottom cover. Success! An opening tool takes care of the base, and six quick stabs with the TR6 Torx security driver loosens the familiar antenna plate underneath. So far so good.

- We're greeted by a single fan standing watch over the insides. The fan unscrews with zero fuss, giving a better view of the mini's depths. A firm push on either side of the blower exhaust, and the board unclips and slides out.

- With the board out, we're going straight for the RAM. Apple has trapped it in a heavy metal cage—almost as if it doesn't fully trust modular RAM to behave itself. Actually, we've seen this in iMacs of yore. The shield allows the RAM to operate at high frequency (2,666MT/s) with no chance of interfering with other functions. Twirl away four Torx screws, and the cage slides off. Has RAM replacement ever been easier? Sure it has—but the return to standard SO-DIMM RAM after the bitter disappointment of the 2014 mini's soldered-down chips is a huge win.

- Just one connector and two screws set the little system speaker free. iMac and MacBook speakers seem to get bigger all the time, but this looks about the same size as in older Mac minis. Beneath the speaker, we find some antenna cables, but no modular AirPort card—in what is becoming a trend, these are socketed to the main board.

- From here, we set to work freeing the heatsink, twirling away Torx screws, and exposing the paste-y (soldered) CPU. One last screw, and the port cover is free, uncovering the ports. As it departs, it takes some antenna hardware with it. The last thing between us and an empty mini is the internal power supply, which is a nice enclosed unit, making for safe, easy replacement.

- Repairability score: 6 out of 10 (10 is easiest to repair). No tough adhesive holds the Mac mini hostage, and disassembly is straightforward. The mini packs standard SO-DIMMs, allowing both DIY upgrades and replacements. The CPU and storage are soldered to the logic board, and not user-upgradeable. If any of the many ports is damaged or worn, the entire logic board will need replacing.

Program and Set Up a Tiny Web Server

YOU'LL NEED THIS

A LINUX PC
This needs to run OpenVPN.

A WEB SERVER is a hefty collection of software running in a rack server, which—unsurprisingly—few of us have at home. So, when you want to make a new webpage, you have to use an online service, such as Codepen or Jsfiddle, to test your CSS code.

However, these services are usually limited to handling small code snippets, and more advanced code becomes cumbersome to develop using such tools. With greater complexity come other risks, as well, such as compatibility. If you hope to be a professional coder, you'll need to know how to use a few tools to ensure that you're building on the correct platform for which you will eventually be deploying. In this article, you'll be able to learn a few techniques to accomplish this. –**MATS TAGE AXELSSON**

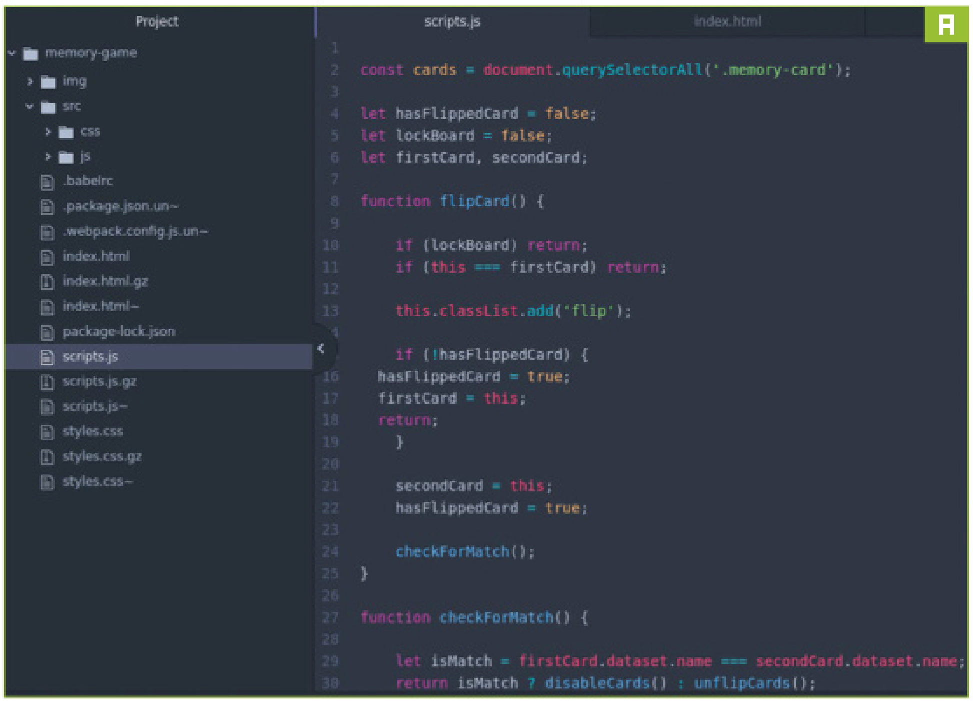

```
const cards = document.querySelectorAll('.memory-card');

let hasFlippedCard = false;
let lockBoard = false;
let firstCard, secondCard;

function flipCard() {

    if (lockBoard) return;
    if (this === firstCard) return;

    this.classList.add('flip');

    if (!hasFlippedCard) {
  hasFlippedCard = true;
  firstCard = this;
  return;
    }

    secondCard = this;
    hasFlippedCard = true;

    checkForMatch();
}

function checkForMatch() {

    let isMatch = firstCard.dataset.name === secondCard.dataset.name;
    return isMatch ? disableCards() : unflipCards();
```

1 THE BASICS

In its early days, most JavaScript was embedded in the HTML code. This was handy, because there were only a few lines of it on each page. However, now that many users have created frameworks from their initial code, separate program files have become necessary. When you start programming in JavaScript, it makes sense to start with separate JavaScript files, and make a clear structure for all your content [**Image A**], including HTML, CSS, and images. This ensures that you know how to include the frameworks that are essential for going beyond the well-known "Hello World" level of programming.

›› To get started, make the typical "Hello World," and put it in a JavaScript file called "app.js."

```
const http = require('http');

const hostname = '127.0.0.1';
const port = 3000;

const server = http.createServer((req, res) => {
  res.statusCode = 200;
  res.setHeader('Content-Type', 'text/plain');
  res.end('Hello World\n');
});

server.listen(port, hostname, () => {
  console.log(`Server running at
http://${hostname}:${port}/`);
});
```

›› At the top of the listing, the code starts a web server and sets the minimal parameters to run it locally. This is the fastest way to get anything running as a local server. Further down, you can see code that runs a "Hello World" script. To start this JavaScript, just run `node app.js`. This is fine if you just want to see how a certain piece of code works. However, when you're involved with bigger projects, there are other solutions that are easier to handle, although they require a module. These modules start a server so that all files in your tree are read as they will be in the live server. If you want help with JavaScript, the Mozilla Foundation has a fantastic set of tutorials on its documentation webpage [**Image B**].

2 SMALL AND PERFECTLY FORMED

Why so many modules in the first place? The answer is that you want the module to be small, yet still support all the functions you may need. You can find an example in the "react-native-httpd" module. Not everyone wants to use React, but you need to have support for the correct frameworks, otherwise the code simply won't work. On the other hand, if you have elements you don't use, the results may differ when you deploy. The main reason is that we follow the classic Unix programming rule of thumb: Make each software tool do one thing and do it well.

›› One of the most common frameworks for JavaScript code is Node.js. To install these packages, you'd usually use the packaged version. This can be a mistake, though, because the newest version isn't usually available there. You may also need to have several versions because of the site's software levels. You're better off going to https://nodejs.org/en/download/package-manager/#debian-and-ubuntu-based-linux-distributions, and finding the package there.

›› The setup script changes your system's list of repositories, so you can keep using the newest version. You can register on www.npmjs.com, but this is only useful when you start sharing code.

›› In the npm registry, you'll find several packages for running a small, local web server. Some are for static content only, while others are only for serving files. The package called "httpd-server" can handle what most regular servers can, and better still, it's relatively small and customizable.

```
$ sudo npm install http-server -g
```

›› Start your server with the command `http-server`. Without parameters, the server starts with the current directory as the main one. This makes it easy to change into your source repository, and run whatever you've created.

›› You can also start the command with parameters pointing to the correct directory. Your source code needs to be organized so you can easily deploy with the same set of files on a live server. In the case of the "http-server," you now have this solution ready. Remember that any files below the directory you start from are invisible to the server.

3

CONFIGURE HTTP-SERVER

To make everything work as closely as possible to the real thing, you can trim some parameters. The port defaults to 8080, but you can pick your own port in case you have other local servers. Parameter `-g` picks a gzip version of any file that you've packed and had the same name before the "gz" extension. You can also change the default HTML extension, which is useful if you're making a PHP application.

›› When you need to use secure webpages, you can also start the use of SSL and point to a certificate and key file. Without the two files, the site won't start, as encryption can't occur—something that people who work with secure sites will already know. To make sure your code always updates in the browser, you can also turn off caching by setting the cache time to -1.

›› What can't you do? Well, because the server needs to be small, you won't have full support for any framework beyond Node.js. The package can also use a proxy function. This sends some requests over to a webpage, even though your code and data is local. With this in place, you can test how to collect information from other sites, such as embedding remote content. You can rectify this by adding it afterward, but before you do that, look through the other options that are available:

```
$ http-server -p 8080 --gzip -c-1
```

›› When using this tiny server, you must compress all files before the server needs them, if you want to use compression. In a more advanced web server, for example nginx, you can choose to compress on demand. Either way, it's useful to be able to measure speed gains of compressing your files.

ONLINE TOOLS FOR TESTING CODE

Before you start setting all this up on your home system, and if you want to play with small snippets of code, try online solutions, such as Codepen and Jsfiddle.

Codepen has a pleasing design and it's easy to move around the desktop. For completeness, you have a window for HTML, one for CSS, and finally one for JavaScript. Each window also has options to analyze your code in different ways. The features are Tidy, Analyze, and Show Compiled. These are useful if you have a bad habit of indenting poorly, or just can't figure out what you've done wrong. You can choose where you want your code to show. The three windows with code can be on the right, left, or bottom. When you want to start your own project, you have several templates available to get started. Your finished project can be downloaded as a zip archive.

Jsfiddle is similar to Codepen, but has more view options and templates. It also has a neat drop-down list where you can choose language, library, and other values for your code. In Jsfiddle, you also have a collaborate function. When you activate this feature, you create a link that you share with someone you trust, and you can both work on the same code together.

These tools are great for testing new ideas and building the initial code. However, if you already have code on GitHub, then you can link directly to it on Jsfiddle.

›› Of course, this solution is one of many. In this section, you'll learn about the possibilities and limitations of this particular one. Other options are web servers built with more flexibility and/or more features. The "local-web-sever" package packs a wider array of features and customizability into it. To install it, use the npm manager:

B

Hello World

At Mozilla, we're a global community of

- technologists
- thinkers
- builders

working together to keep the Internet alive and accessible, so people worldwide can be informed contributors and creators of the Web. We

```
$ npm install -g local-web-server.
```

›› Once that's done, the simplest case is to host a static site with no arguments. This is the same as the former package. The command is `ws`. The default port is 8000, which you can choose yourself. If you want to serve a Single Page Application, you use the argument `--spa`. You can also use proxies in this package.

›› To serve https and http2, start with the arguments `--https` or `--http2`, respectively. There are built-in keys and certificates, so you can run without your own keys, although you can also supply your own. In case you're interacting with another site, be it on the web or another part of your own, you can set up mock answers. These mock answers can send a reply to a request from an existing site using a preset answer. This way, you can test when you have no contact to the specific site, even when you are offline. To define mock answers, use a JavaScript file like the one below.

```
module.exports = MockBase => class MockPaths
extends MockBase{
  mocks () {
    return  {
      route: '/paths',
      responses: [
        {
          response: {
            type: 'json',
            body: [
                    {name : 'Roslagsleden',  goesThrough =
'Roslagen'},
                    {name : 'Sörmlandsleden',  goesThrough
= 'Sörmland'},
                    {name : 'Nordkalottleden',  goesThrough
= 'Norway'},
                  ]
                  }
        }
      }
    }
}
```

» Save the code above, call it "mock-example.js," then run the web server.

```
$ ws --mocks mock-example.js
```

» After this, you can get the response with `curl`.

```
$ curl http://127.0.0.1/paths
```

» Now you can request your data from the local system to check that your code can handle the response.

» As your projects grow, your testing requirements also grow. To help verify functionality in many scenarios, you can use the "dev-web-server," which has more functions. It also supports defining API endpoints—and with these functions, you can call APIs. You also have the chance to define both success and failure of the endpoints. With all this functionality, the full design and the testing can be done locally, making the whole flow safer and less prone to delays that lie beyond your control.

» The "dev-web-server" package also has a feature for failure scenarios. To set everything up, consider the below command to start the server:

```
$ dev-web-server DOMAIN 0.0.0.0 PORT 1234 BASEDIR
..\rep\httpdocs DELAY 2000 ENDPOINTS ..\rep\server\
my-endpoints.js
```

» As you can see, the delay is set to 2000. Using this argument, you can test how things will work when your user has a big lag to the webpage. All parameters can also be defined by the "dev-web-server.json" file in the launching directory. Any arguments on the command line override the file. The example file defines paths and delay. It also defines whether you use CORS or not. CORS is a way to enable programs to fetch data from other domains that they co-operate with. This is regularly blocked because of security concerns.

```
{
  "domain": "0.0.0.0",
  "port": 13002,
  "baseDir": "./test/test-pages",
  "delay": 0,
  "endPointsFilePath": "./api-proxy/api.js",
  "withCORS": "true"
}
```

» The other parameters set the base directory where you put your webpage files, and finally the endpoint file. In the endpoint file, you define how responses are handled. These responses include successes as well as failures. The format is a JavaScript file where functions are defined. The example points to the "api.js" file, and part of that file is shown below.

```
var Repository = {
  count:0
};
 module.exports = {
   '/example': function (req, res, params, sendSuccess,
sendError) {
     // Response result is: '{"test":"Cool,
success","count":1}'.
     // HTTP code is 200
     sendSuccess(req, res, {
       test: 'Cool, success',
       count: Repository.count++
     });
   },
   '/exampleError': function (req, res, params,
sendSuccess, sendError) {
     // Response result is: '{"code":401,"message":"You
need to login"}'.
     // HTTP code is 401
     sendError(req, res, 401, 'You need to login');
   }
};
```

» In reality, you'll cover many more error codes, and these are two of the main ones. In addition, note that the "dev-web-server" package has two dependencies: colors and object-tools. This is not usually a problem, but worth mentioning anyway.

4

LIVE RELOAD FACILITY

If you use the "webpack" module, you'll be interested in "webpack-dev-server" [**Image C**]. As the name suggests, the package is designed for development only. One of the most interesting features is the live reload feature, which makes all your new changes load immediately. Often, when you find a problem in your code, changing the code is done quickly, but then you have to make sure that the code is actually used by the server. The browser and server cache go some way toward saving bandwidth for users, but when you're coding, this isn't always obvious, while what level your browser is using isn't always clear.

» For the uninitiated, the "webpack" module is designed to take your project and its dependencies, and pack them into a single module. This means that anyone interested in your module can download the pack and use it as is, without worrying about dependencies. In the case of "webpack-dev-server," the pack remains in the memory and is served to the client.

» For npm projects to work at all, you must initialize your directory. This creates the "node_modules" directory that contains all the modules you'll use.

```
$ npm init -y
```

» Without the `init`, packages won't be installed correctly to the directory; you can install globally, but you have greater control if you make an install local.

» The package manager will handle dependencies for you, but you need to tell it what you're after. You need to make sure the "webpack" module is available on your system. In case it isn't, install it with:

```
$ npm install webpack --save-dev
$ npm install webpack-cli --save-dev
```

» The `save-dev` argument means that the package will only be used for development. This is important, because you may need some packages for developing, while some are critical on the deployed server. You'll get some warnings, but you can safely ignore the "fsevents" dependency warning as it's not supported on Linux.

» Next, you need to adapt the configuration file for your own needs. In the file, your main concern is the

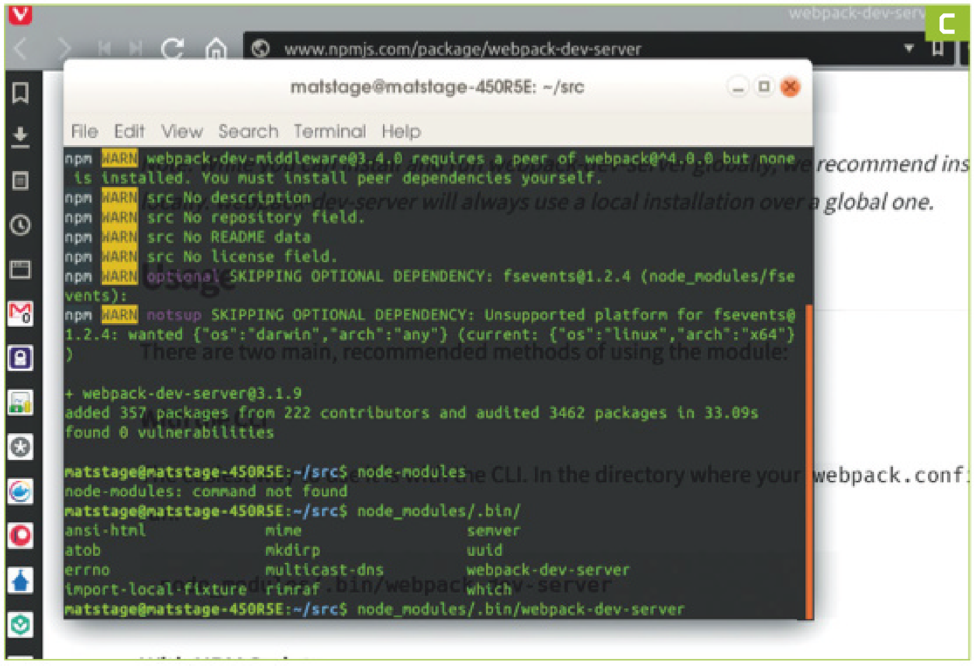
www.npmjs.com/package/webpack-dev-server
matstage@matstage-450R5E: ~/src
File Edit View Search Terminal Help
C

```
npm WARN webpack-dev-middleware@3.4.0 requires a peer of webpack@^4.0.0 but none
 is installed. You must install peer dependencies yourself.
npm WARN src No description
npm WARN src No repository field.
npm WARN src No README data
npm WARN src No license field.
npm WARN optional SKIPPING OPTIONAL DEPENDENCY: fsevents@1.2.4 (node_modules/fse
vents):
npm WARN notsup SKIPPING OPTIONAL DEPENDENCY: Unsupported platform for fsevents@
1.2.4: wanted {"os":"darwin","arch":"any"} (current: {"os":"linux","arch":"x64"}
)

+ webpack-dev-server@3.1.9
added 357 packages from 222 contributors and audited 3462 packages in 33.09s
found 0 vulnerabilities

matstage@matstage-450R5E:~/src$ node-modules
node-modules: command not found
matstage@matstage-450R5E:~/src$ node_modules/.bin/
ansi-html           mime                semver
atob                mkdirp              uuid
errno               multicast-dns       webpack-dev-server
import-local-fixture rimraf             which
matstage@matstage-450R5E:~/src$ node_modules/.bin/webpack-dev-server
```

entry point. This is the JavaScript file where "webpack" starts to look through your application. In the case of a simple application, you may only have one JavaScript file, but when it grows in size, your best bet is to have many. A great example is available on GitHub (https://github.com/mschwarzmueller/yt-webpack2-basics). This example is focused on "webpack," but includes some details about the web development server. Initially, you need to create your "package.json" file in which you define how the package is run:

```
{
  "name": "webpack2-basics",
  "version": "1.0.0",
  "description": "",
  "main": "index.js",
  "scripts": {
    "test": "echo \"Error: no test specified\" && exit 1",
    "run-server": "webpack-dev-server --entry ./src/js/
 app.js --output-filename ./dist/bundle.js",
    "build:prod": "webpack src/js/app.js dist/bundle.js
 -p"
  },
  "author": "",
  "license": "MIT",
  "devDependencies": {
    "webpack": "^2.2.1",
    "webpack-dev-server": "^2.4.1"
  }
}
```

›› In the directory of the "package.json file," run `npm`. The `run-server` argument is taken from the "package.json" setting. In this case, you can use `build:prod` to make a bundle file:

```
$ npm run run-server
```

›› You should now be able to see the site in your web browser. The port will show up in your terminal window.

›› For more advanced use of "webpack" use the "webpack.config.js" file. Creating it may be a little awkward in the beginning. You can always try the tool at https://webpack.jakoblind.no, which creates the configuration files for you. The result will be a good starting point.

5 PARALLEL EXECUTION

At first, your webpages will be trivial, but it's still good to know where to look for the packages that can make your system run in parallel. There are several solutions to this; some known ones are "forky" and "throng." Of the two, "forky" is the more advanced, but you need to activate both in a JavaScript function. "Forky" has one master file and several process files, while "throng" has only worker files and processes. Whichever you choose, it's good practice to consider these issues before you think you'll need them.

THE REACT FRAMEWORK

React is developed to handle user interfaces, so you should think of it as a part of your toolset, not the entire solution.

To use React, you may want to consider what logic you need to combine with it. Combining React with a CMS is one way to go, and there's even a webpage (https://reactfordrupal.com) for this purpose. Drupal is not the only choice, though. WordPress has an API that makes developing the interface straightforward.

When you have npm installed, you can look for packages using the built-in search. The result is a long list of packages, so be sure to choose "react-dom" for webpages; "react-native" is for native environments. Packages that use React make an even longer list. In fact, the length alone should encourage you to learn React and consider using it for your next project. If you think about how much it's used in live installs, you'll be even more motivated. React is a major force in the move toward a headless operation that most CMSes are moving to, and if you don't want to use a CMS, you can still use it for your regular application. It's perfect for one-page applications, where the page behaves more like an application than a webpage.

›› For the more professionally minded, you should learn how to build a containerized environment, so the correct system can be available to you, irrespective of where you'll finally deploy your code. To make absolutely sure, you need to build an image of the site you're deploying to. At first, this will seem as though you end up with the sort of gigantic server that this whole exercise was designed to avoid, but it's not quite like that.

›› If you use a containment system such as Docker (www.docker.com), you can have the whole thing running without touching your system, and the "web server" you'll have created will act just like the real one. A good starting point for web servers is to use a minimal image—the smaller ones are just a few hundred kilobytes. Even the Ubuntu one with nginx on top is quite small, but managing packages is another issue when using Docker. That is a story for another time…

6 MOVING ON…

In this tutorial, you've learned the basic methods for starting a web server on your local machine for testing. A tiny specialized server won't overload your local hardware—a simple laptop is enough. Npm has a number of great solutions for this type of work. For bigger jobs, you'll most likely need to use container technology like a Docker image to do the same. For your early learning and experimenting sessions, though, npm will suit you down to the ground. ⏻

Build a Pi-Powered Random Video Player

YOU'LL NEED THIS

ANY MODEL OF RASPBERRY PI
Plus the latest version of Raspbian; three push buttons; breadboard; four female-to-male jumper wires; two male-to-male jumper wires.

STUCK ON WHAT TO WATCH this evening with the family? Fear not—with this project, we can have our Raspberry Pi randomly choose a video for us to watch. So, let's start building and coding our own random video player. All of the code, and a detailed high-res circuit diagram, can be downloaded from https://github.com/lesp/LXF244-VideoPlayer/archive/master.zip.

Building the circuit for this project is simple. We only require three buttons, connected to the GPIO at pins two, three, and four, and each of these buttons also needs to be connected to GND. For this we can use a single GND pin on the Pi, and a female-to-male jumper wire. Connecting to the negative rail of a breadboard means that we have multiple GND connections that can be connected via the two male-to-male jumper wires. See the diagram in the download for this project for details. When the hardware is built, attach the accessories required for your Pi, then boot to the Raspbian desktop.–**LES POUNDER**

1 TUNE IN

Before we can write any Python code, we need to install two libraries that will help us create the project. Open a terminal, and type the following:

```
$ sudo pip3 install glob
$ sudo pip3 install keyboard
```

›› We'll talk more about them later. For now, open the Python 3 editor, found in the "Programming" menu, and click "File → New" to create a new blank file. Immediately click "File → Save" in the new window. Save the code as "VideoPlayer.py" and remember to save often.

›› Now we start writing the code [Image A], and our first act is to import the libraries that we require. The first three libraries are used to detect the button presses (GPIO Zero), stop the code from running once and exiting (pause), and to choose a random video (choice.)

```
from gpiozero import Button
from signal import pause
from random import choice
```

›› The final three imports are a library that we shall use to list the contents of a directory (glob), one to run terminal commands (subprocess), and the last is a library to emulate a keyboard (keyboard).

```
import glob
import subprocess
import keyboard
```

›› Moving on, we create three functions—blocks of code that we can later reuse by calling their name. The first function is called "play_video," and it first creates a list (a data storage object) called "videos."

```
def play_video():
    videos = []
```

›› To fill the list with all the videos we can play, we use a "for" loop. This will iterate over every file in a directory, as long as it's an MP4 video file. For this we use `glob` to access the directory "/media/pi/Videos," which is really a USB stick called "Videos" full of MP4 files. Change this to match the name of your chosen directory full of videos. Each time an MP4 is found, it's appended to the "videos" list that we've just created.

```
or file in glob.glob("/media/pi/Videos/*.mp4"):
    videos.append(file)
```

›› Check that the list has been populated with file names by printing the list's contents to the Python shell:

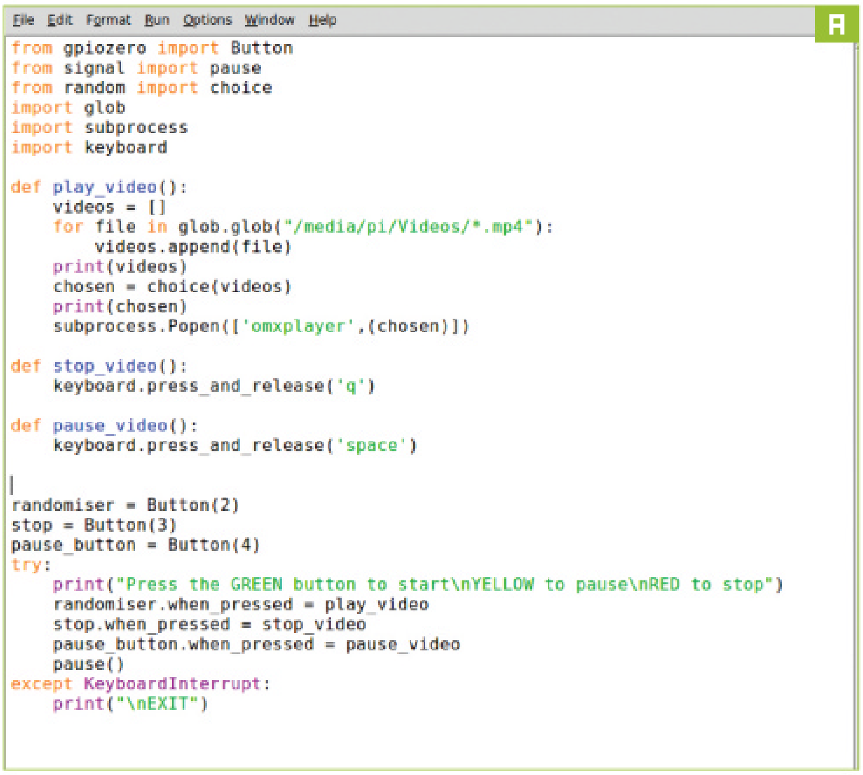

```
from gpiozero import Button
from signal import pause
from random import choice
import glob
import subprocess
import keyboard

def play_video():
    videos = []
    for file in glob.glob("/media/pi/Videos/*.mp4"):
        videos.append(file)
    print(videos)
    chosen = choice(videos)
    print(chosen)
    subprocess.Popen(['omxplayer',(chosen)])

def stop_video():
    keyboard.press_and_release('q')

def pause_video():
    keyboard.press_and_release('space')


randomiser = Button(2)
stop = Button(3)
pause_button = Button(4)
try:
    print("Press the GREEN button to start\nYELLOW to pause\nRED to stop")
    randomiser.when_pressed = play_video
    stop.when_pressed = stop_video
    pause_button.when_pressed = pause_video
    pause()
except KeyboardInterrupt:
    print("\nEXIT")
```

```
    print(videos)
```

Then choose a random video from the list and store it in a variable called "chosen." Print this to the Python shell:

```
    chosen = choice(videos)
    print(chosen)
```

›› The last line in this function uses the `Popen` class from the subprocess library to run a command as if we were sitting at the terminal, and this command is to open the omxplayer media player, then play the chosen video.

```
    subprocess.Popen(['omxplayer',(chosen)])
```

›› The next function is called "stop_video," and as you can guess, it stops the currently playing video. For this we use the keyboard library, specifically the "press_and_release" function, to simulate pressing Q.

```
def stop_video():
    keyboard.press_and_release('q')
```

›› The last function is called "pause_video," and it emulates pressing the space bar on the keyboard, which is how omxplayer pauses video.

```
def pause_video():
    keyboard.press_and_release('space')
```

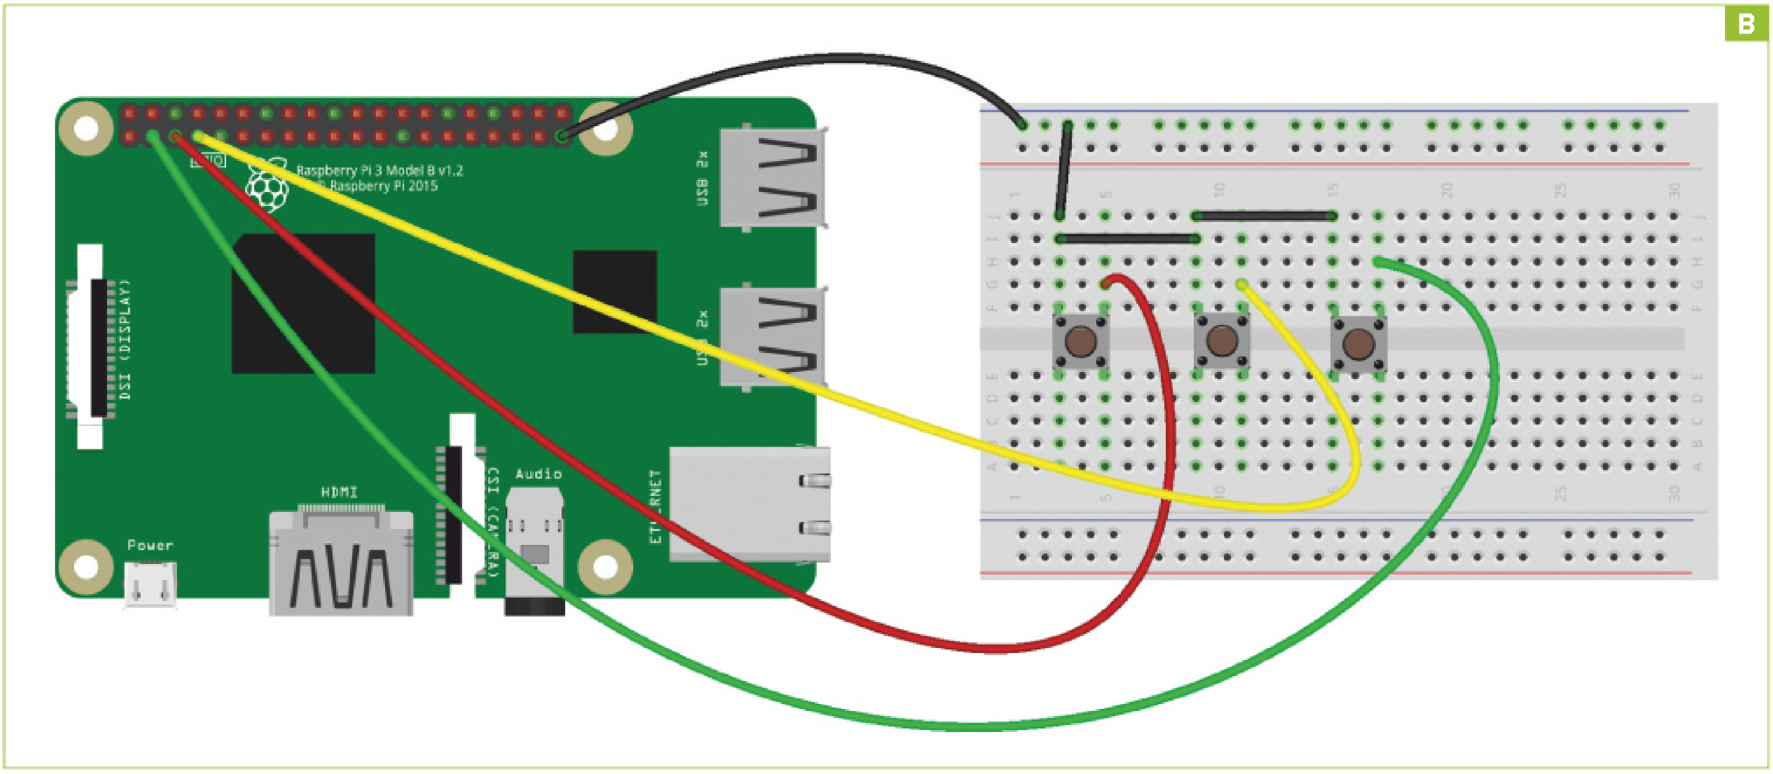

» With the functions created, we next need to tell our code where our buttons are connected. We have three buttons connected: randomiser (play), stop, and pause_button. These buttons are connected to the GPIO at pins two, three, and four, respectively [**Image B**].

```
randomiser = Button(2)
stop = Button(3)
pause_button = Button(4)
```

2

PUSH THE BUTTON

On to the last part of the code, which looks for a video when a button is pressed, and then reacts accordingly. First, however, we need to wrap this section in an exception handler. It tries to run the code, but if the user presses Ctrl-C, the code exits. So, for the `try` section, our code first prints three lines to the shell—these are the instructions to the user. The `\n` between each color is Python shorthand to insert a new line between each instruction:

```
try:
   print("Press the GREEN button to start\nYELLOW to
  pause\nRED to stop")
```

» Our three buttons are waiting to be pressed, and using the GPIO Zero library, we call the `when_pressed` class to detect when each of the buttons are pressed. When this happens, the corresponding function is run. So, when we press the green randomizer (play) button, it randomly chooses a video from the USB stick, and plays it for us. You'll notice that the functions don't have a `()` at the end of the function name. If they did, the code would run automatically. Right now, they're ready to run on demand.

```
   randomiser.when_pressed = play_video
   stop.when_pressed = stop_video
   pause_button.when_pressed = pause_video
```

» The last line in this section is a simple `pause()`, and this is used to keep the code running, and not just exiting after running once.

```
   pause()
```

» But what happens if the user presses Ctrl-C? Well, this is a "Keyboard Interrupt," and if that occurs, we shall simply print an empty line, then "Exit" on the screen.

```
except KeyboardInterrupt:
   print("\nEXIT")
```

» Now save your code, and open a terminal to the directory where you have saved the code. In order to use the code, we need to run it with `sudo` (root powers), because the keyboard library used in the project can only be used as root. To run the code, type:

```
$ sudo python3 VideoPlayer.py
```

When you're all sitting on the couch, press the green button to play the randomly chosen video, to pause press yellow, and press red to stop. Enjoy your evening, and be careful not to spill the popcorn!

KEYBOARD SHENANIGANS

In this project, we used the Keyboard Python library to emulate a real keyboard using just the push buttons that controlled playing, pausing, and stopping video playback. The Keyboard library can be used to press any keys, so we can automate multiple key presses all from one push button. The library can also be used to write text to the screen.

However, the most dangerous function of this library is that it can be used to record every single key press on the target machine. These key presses can be recorded to a list, then saved to a file, and used for nefarious purposes. Of course, recording the key presses of a user, without their consent, is illegal, and can land you in a lot of trouble—so don't do it! With that said, it can be used as a powerful tool when debugging how a user interacts with your code, so for software testing (under consent with the general public and in-house testers), you can see which keys they were pressing right before the code locked up and/or went thermonuclear!

The Keyboard library works with Linux and Windows, and it's certainly lots of fun to play with. To read more about this library, head over to the blog post at https://bigl.es/tuesday-tooling-record-replay-keystrokes-with-python, and see a few examples of how it can be used in day-to-day life.

Fix OneDrive Sync Problems

YOU'LL NEED THIS

WINDOWS & ONEDRIVE
OneDrive offers 5GB of free storage, but there's a number of paid-for plans available if you need more.

WHEN IT WORKS, Microsoft's OneDrive service is fantastic. It's built into Windows 10, and enables you to easily synchronize files and folders with other Internet-connected computers, tablets, and smartphones. It means your files are easily accessible from a range of devices—provided they can get online. However, OneDrive can occasionally encounter problems synching your files. This can be frustrating, especially if you're away from the PC the files originated from, because your files won't be properly synchronized with your other devices, and that can make them impossible to download, edit, or share.

The good news is that these issues are relatively rare, and if you do encounter them, they're quite straightforward to fix. So, before you start panicking that your files aren't synching, follow the steps below, and we'll have your cloud-computing issues fixed in a jiffy. **–MATT HANSON**

1 CHECK ONEDRIVE IS RUNNING

Make sure the OneDrive app is running in Windows, and you've signed in correctly. When you sign up for a Microsoft account—which is advised if you want to log in to Windows 10, personalize it, and download apps from the Microsoft Store—you also get OneDrive, which should load when Windows starts. If not, search for "OneDrive" from the Taskbar [**Image A**].

2 CHECK THE ICON

Keep an eye on the OneDrive icon (which looks like two clouds) in the Taskbar's notification area. This gives you a quick overview of how OneDrive is synching, or if there's an error. If the icon appears normal, it means your files are synched. If it has a small icon of two arrows in a circle [**Image B**], it's synching. If there's a red cross, there's a problem.

3 FIND OUT MORE

If you see the red cross icon, click it, then select "View sync problems." This opens up a dialog box, which gives you more information about the error. In this window, you'll see the name of the file that's causing the issue. You'll also see an "Error" column. Here, you'll get details about the problem, which should hopefully give you an idea of what's going wrong.

4 CHECK YOUR INTERNET CONNECTION

The previous steps solve most issues, but what if you don't get the red cross icon, yet there's still a problem? First, check you're connected to the Internet—OneDrive needs access for it to work, so if your PC isn't connected via a network cable or Wi-Fi, OneDrive fails. In the notification area of the Taskbar, you'll see an icon that shows if you're connected to the Internet.

5 DOWNLOAD THE LATEST VERSION

If you've connected to the Internet without issue, make sure you have the latest version of OneDrive installed. In Windows 10, as long as you've kept the OS up to date, you should have the latest version of OneDrive, too. Type "Windows Update" in the search bar, and make sure you're updated [**Image C**]. You can download the latest version from http://bit.ly/LatestOneDrive.

6 UNLINK AND RELINK ONEDRIVE

Unlinking OneDrive, then relinking to it can prompt it to sync properly again—don't worry, you won't lose any files or folders you've already synched. To do this, click the OneDrive icon in the notification area, select "More" in the window that appears, then click "Settings" [**Image D**]. In the next window, select "Unlink this PC," and follow the instructions to relink the PC to OneDrive.

7 CHECK YOUR STORAGE SPACE

While you have the "Settings" window open, it's a good idea to check how much space you have left on OneDrive. You are given a set amount of storage space, and if you run out, you won't be able to sync any more files. You can either buy more space, or click "Manage Storage" [**Image E**] to find the files that are taking up the most space, and remove them from your OneDrive storage.

8 REMOVE PROBLEM FILES THEN TRY AGAIN

You can also try moving the files that are attempting to sync out of the OneDrive folder, and trying again. To do this, click "More" in the OneDrive app, select "Pause synching," and choose any length of time [**Image F**]. Select the file OneDrive is struggling with, and move it to a folder that isn't synched with OneDrive. In the OneDrive app, select "Resume synching." Once complete, move the file back.

9 MAKE SURE FILES CAN BE UPLOADED

If OneDrive is struggling with a file, make sure it's supported. For example, OneDrive is unable to upload files that are larger than 20GB. If it's larger than that, consider putting it in a ZIP file to make it smaller [**Image G**]. Also, make sure the file name isn't too long, because that can confuse OneDrive, too. If the entire file path, including the file name, exceeds 400 characters, it won't upload.

10 RESET ONEDRIVE

Hopefully, one of the previous steps will have sorted out your OneDrive issue. If it hasn't, you may need to reset OneDrive. Again, this doesn't delete any of your files. Press Win-R to open the Run dialog. Now type "%localappdata%\Microsoft\OneDrive\onedrive.exe /reset" [**Image H**]. Press "OK." After a few seconds, the process will complete. You can now relaunch the OneDrive app.

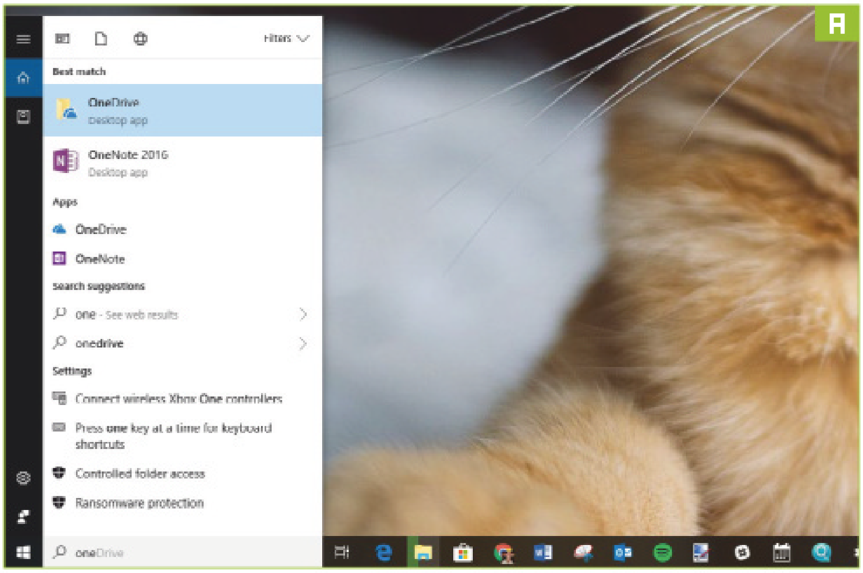
A
Filters
Best match
OneDrive
Desktop app
OneNote 2016
Desktop app
Apps
OneDrive
OneNote
Search suggestions
one - See web results
onedrive
Settings
Connect wireless Xbox One controllers
Press one key at a time for keyboard shortcuts
Controlled folder access
Ransomware protection
oneDrive

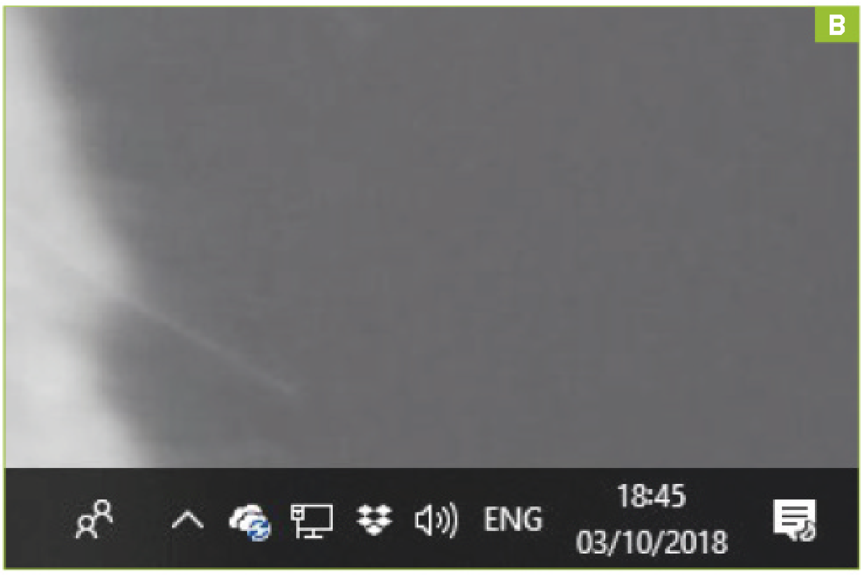
B
ENG
18:45
03/10/2018

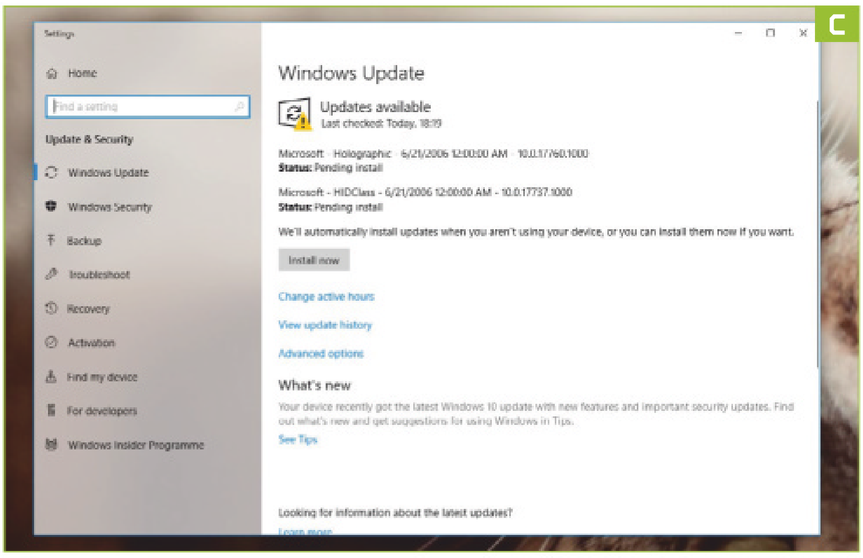
C
Settings
Home
Find a setting
Update & Security
Windows Update
Windows Security
Backup
Troubleshoot
Recovery
Activation
Find my device
For developers
Windows Insider Programme

Windows Update

Updates available
Last checked: Today, 18:19

Microsoft - Holographic - 6/21/2006 12:00:00 AM - 10.0.17760.1000
Status: Pending install

Microsoft - HIDClass - 6/21/2006 12:00:00 AM - 10.0.17737.1000
Status: Pending install

We'll automatically install updates when you aren't using your device, or you can install them now if you want.

Install now
Change active hours
View update history
Advanced options

What's new

Your device recently got the latest Windows 10 update with new features and important security updates. Find out what's new and get suggestions for using Windows in Tips.

Looking for information about the latest updates?

Learn more

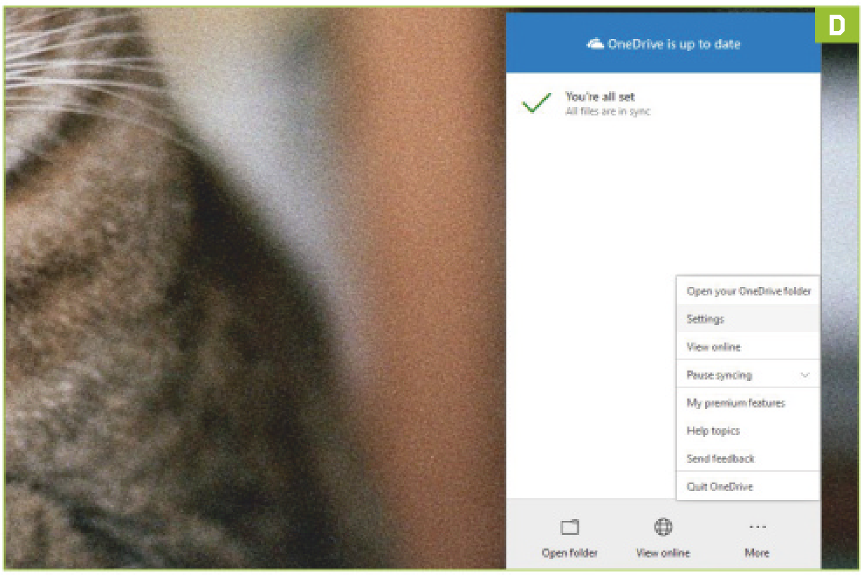
D
OneDrive is up to date

You're all set
All files are in sync

Open your OneDrive folder
Settings
View online
Pause syncing
My premium features
Help topics
Send feedback
Quit OneDrive
Open folder
View online
More

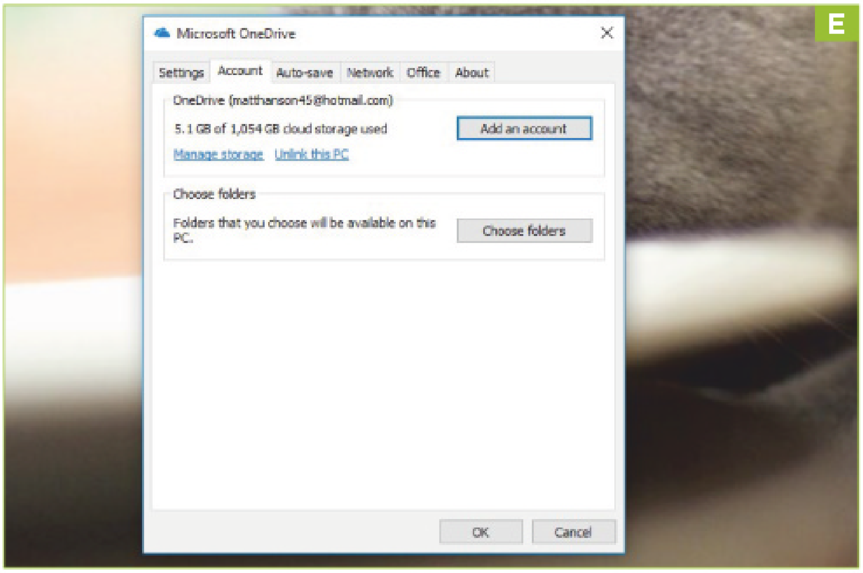
E
Microsoft OneDrive
Settings
Account
Auto-save
Network
Office
About

OneDrive (matthanson45@hotmail.com)

5.1 GB of 1,054 GB cloud storage used

Add an account
Manage storage
Unlink this PC

Choose folders

Folders that you choose will be available on this PC.

Choose folders
OK
Cancel

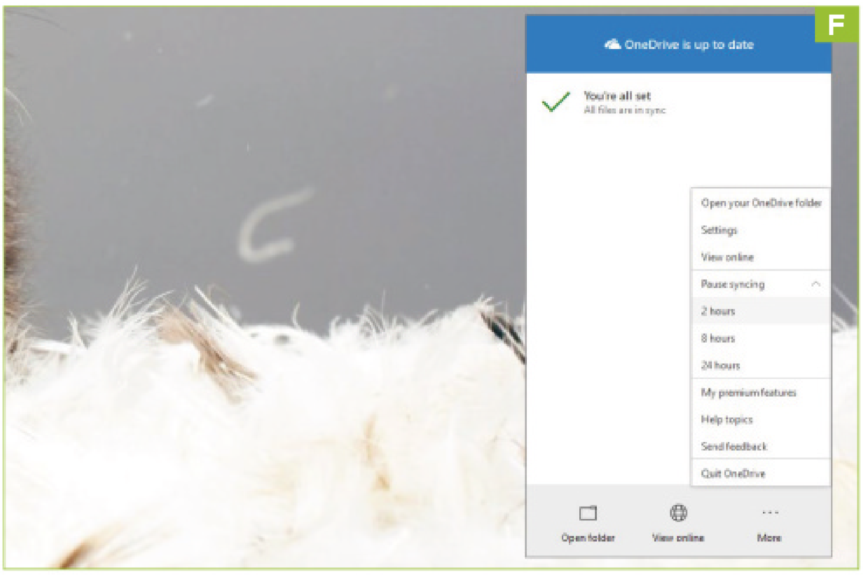
F
OneDrive is up to date

You're all set
All files are in sync

Open your OneDrive folder
Settings
View online
Pause syncing
2 hours
8 hours
24 hours
My premium features
Help topics
Send feedback
Quit OneDrive
Open folder
View online
More

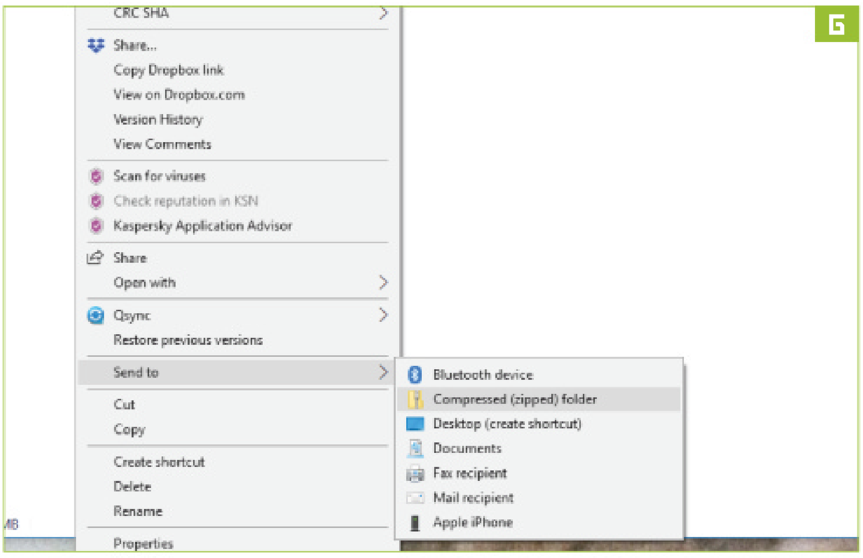
G
CRC SHA
Share...
Copy Dropbox link
View on Dropbox.com
Version History
View Comments
Scan for viruses
Check reputation in KSN
Kaspersky Application Advisor
Share
Open with
Qsync
Restore previous versions
Send to
Cut
Copy
Create shortcut
Delete
Rename
Properties
Bluetooth device
Compressed (zipped) folder
Desktop (create shortcut)
Documents
Fax recipient
Mail recipient
Apple iPhone

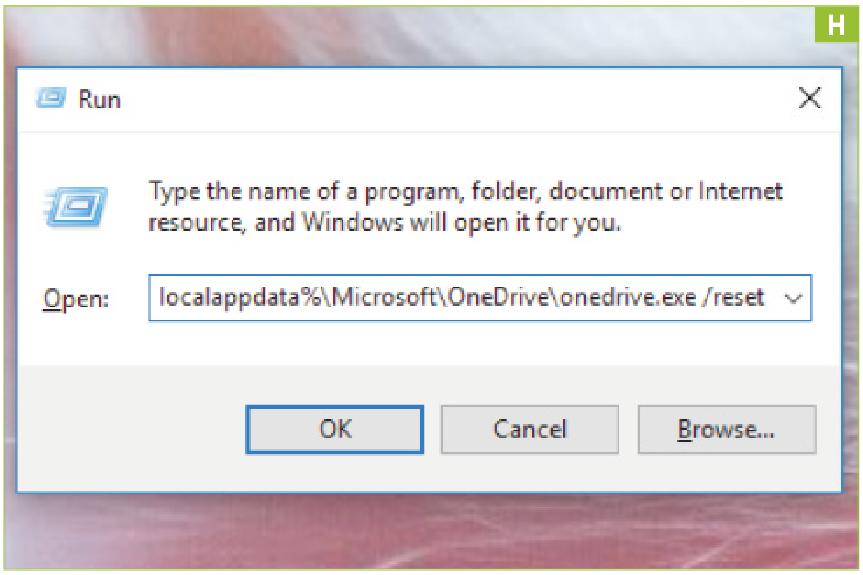
H
Run

Type the name of a program, folder, document or Internet resource, and Windows will open it for you.

Open: localappdata%\Microsoft\OneDrive\onedrive.exe /reset

OK
Cancel
Browse...

Create an Animated GIF with Ease

YOU'LL NEED THIS

ADOBE PHOTOSHOP
You can subscribe to various packages at www.adobe.com.

DESPITE BEING AT THE HEIGHT of its popularity on Geocities back in 1999 (it was the third most visited webpage at the time of its takeover by Yahoo!), the animated GIF refuses to die. Now a staple of the witty reply on Twitter (where the GIFs are converted into more efficient video files, as Twitter doesn't have infinite storage), the GIF seems like it's never going to go away. Even Facebook caught up eventually, adding support for GIFs, originally launched in 1987, in 2015.

Animations came about as an extension to the reworked 1989 GIF specification. Each frame is stored within the same file, and they are displayed in succession with a time delay between. GIFs don't have to loop, and you can set the number of times they repeat—a setting of "0" produces a GIF that loops indefinitely.

There's also a debate over its pronunciation. The creators of the Graphics Interchange Format insist it should be pronounced "Jif," like the peanut butter. However, the first word is clearly "graphics," not "giraffe," so we favor the hard G. It makes no difference to the written word, so pronounce it in your own head any way you like. **–IAN EVENDEN**

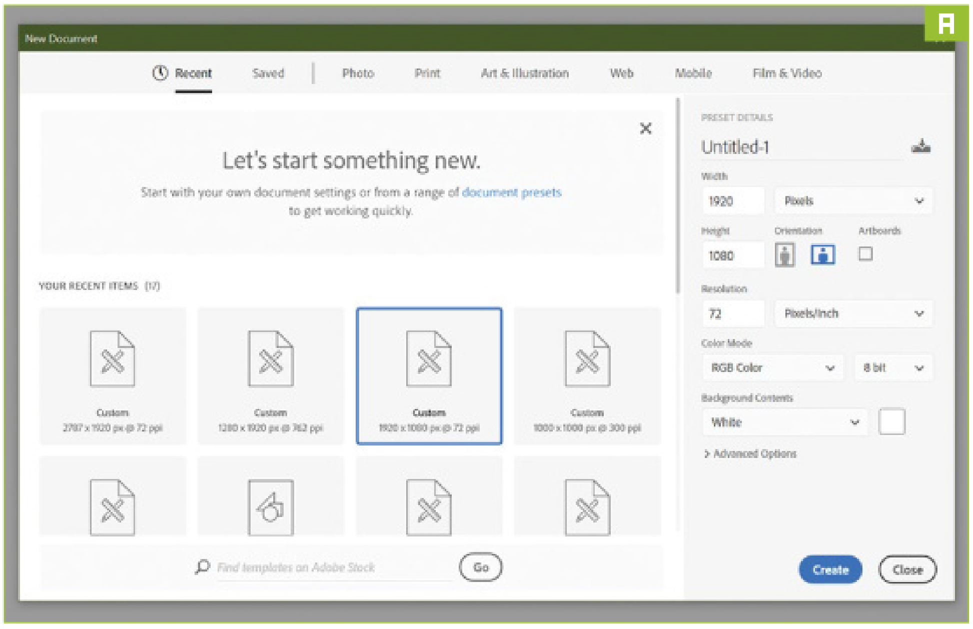

Let's start something new.

Start with your own document settings or from a range of document presets to get working quickly.

Untitled-1

Width: 1920 Pixels

Height: 1080

Resolution: 72 Pixels/Inch

Color Mode: RGB Color 8 bit

Background Contents: White

1 CREATE A BLANK DOCUMENT

It's not often we start with a completely blank document, but here we need to. We've created one at 1920x1080 pixels [**Image A**], mainly so we have something large that will look good in screenshots, but you may choose to make yours smaller in order to create a smaller file at the end. GIFs can display a full range of colors through a block-tiling trick, but the palette table of each block only has 256 entries, meaning 256 colors (commonly 255 plus transparent). This means an animated GIF isn't the ideal format for photorealistic images, especially if you care about such things as quality and file size. What it is good for, however, is graphics, cartoon-like images, and text.

2 CREATE SOMETHING

Simple, eh? Create something. What you create needs to be reasonably simple, so you can draw it again and again, only slightly differently each time as you create the frames of your animation, or something that can be copied and pasted multiple times before being tweaked for each frame. We've created two layers of text at 300pt, and flipped one so it looks a bit like a shadow under the other. We also removed the background layer, using the trashcan button at the bottom of the Layers palette, so we have the checkerboard pattern of a transparent document.

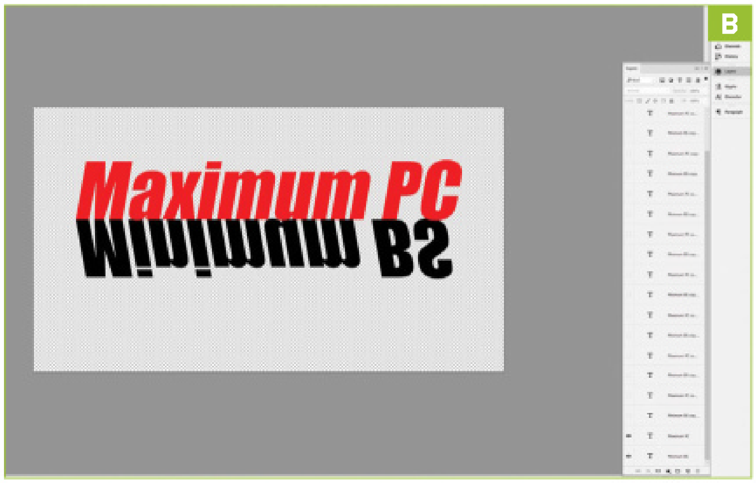

3 CREATE FRAMES

Decide how your image needs to change between the first and last frame, and create the frames by changing your image, with each complete frame having its own layer. We copied and pasted the text we'd created on to 20 layers [**Image B**]—10 for the red text and 10 for the black—and changed each black one a little more using "Edit → Transform → Skew," and its blending opacity on the Layers palette, so that the shadow text appeared to move and strengthen below the red, which never moves. Then we used "Layer → Merge," along with the Layers palette and Ctrl-Left Click to select multiple layers, to create 10 final layers, each with red

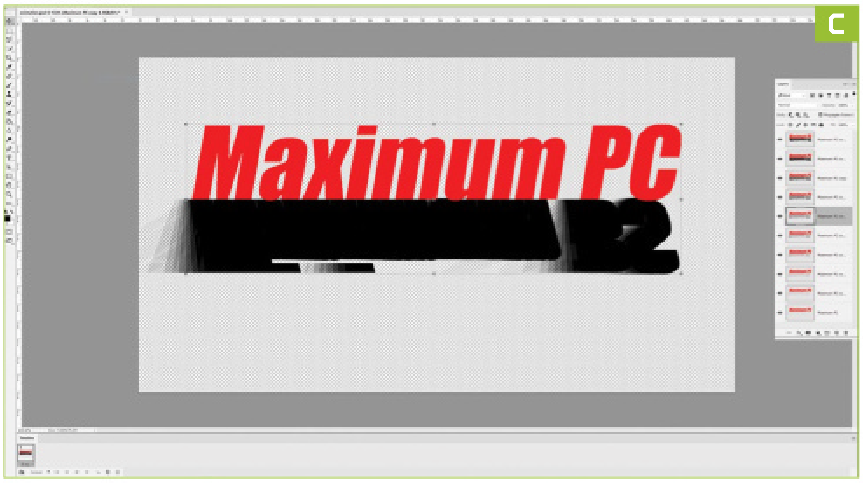

and black text, moving from the beginning at the bottom of the palette to the end at the top. Showing and hiding layers using the eye icon at the left of their entry in the Layers palette can make this easier to manage, and if you're duplicating layers, try linking them by selecting them both, and choosing "Layer → Link Layers" before duplication—this keeps red and black together. Just remember to unlink them before you start editing, or you could get unwanted effects.

4 MERGE FRAMES

Make all your layers visible, then use Ctrl-Left Click to select a red and black pair, before merging them with "Layer → Merge Layers." Once this is done, make sure they're all visible, and open the Timeline from the "Window" menu [**Image C**]. It appears at the bottom of the screen, and at first doesn't have anything on it except a drop-down with two options: "Video" and "Frame animation." Choose "Frame animation" and, in a moment of counter-intuitive design unusual for Adobe, it becomes a button. Click it. The timeline opens, but still looks bare. To fix this, open the three-line menu at its far-right, and choose "Make Frames from Layers." Now you have one frame for each of the layers you prepared earlier. Hit the play button at the bottom-left of the Timeline to preview the final effect.

5 EDIT DELAYS

Each frame in the Timeline has a little drop-down at its base, which currently reads 0 sec. This is the delay on that frame before the next one is shown, and currently there are no delays. You can change this to up to 10 seconds with a click, or enter a custom value [**Image D**]. Beware, though, that it can make your GIF look jerky, and isn't ideal for animation. If you're using the GIF to display a slideshow, however, this is the perfect way to make your images stay on screen longer.

6 SAVING

Once you're happy with how your animation looks, you need to export it. Save the file first as a PSD, in case you want to come back later and edit anything or check your settings. Then head to the "File" menu, and choose "Export → Save for Web (Legacy)." Adobe has been trying to replace this relic of pre-CC Photoshop for a while, but it always has its uses, and this is one of them. It correctly guesses you want to export a GIF, but there are color palette options you can play with to keep the file size down. A "Restrictive" palette produced the smallest file, but watch out for dithering if you've got too many colors in your animation. Setting the number of colors to "Auto" also pares down the table for a smaller file, but the best way to shrink it down is to reduce the resolution, even if you created it at a higher one. You can do this from the "Save for Web" window by choosing new pixel dimensions at the bottom-right, where you can also choose how many times the GIF will loop, or whether it will go on forever. A quirk of this window is that hitting the "Done" button doesn't mean your image gets saved—you need to specifically press the "Save" button [**Image E**]. ⏻

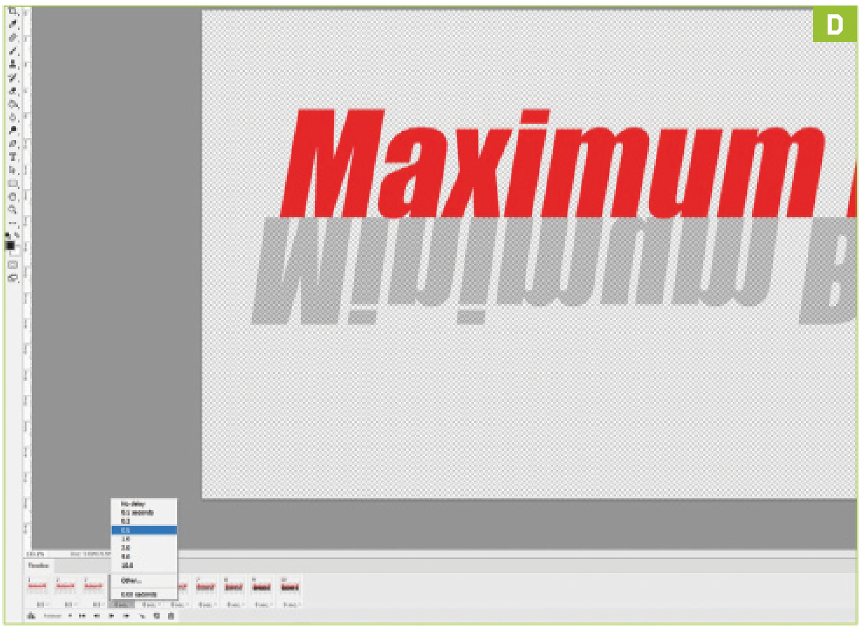
D

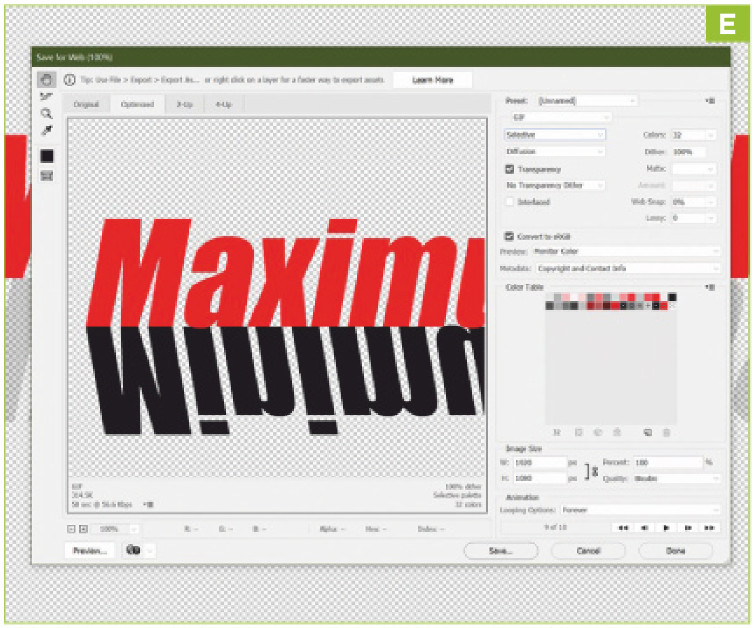
E

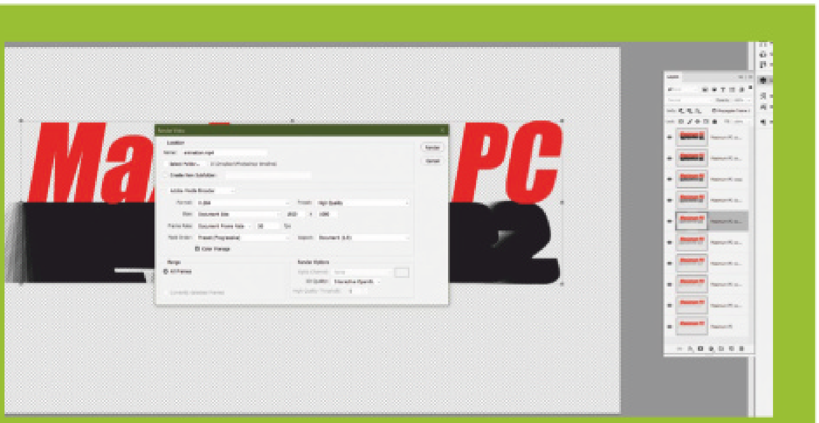

WHAT IF I DON'T LIKE GIF?

Your animation doesn't have to be tied to GIF. Head down the "File" menu, and you'll find "Render Video" nestled under "Export." This enables you to save as an MPEG4 video file, and there are a few options to lower the file size, including changing the resolution. One benefit of video is that it avoids the 256-color limit baked into GIF.

Alternatively, there are extensions to other image formats that can help, if you can find a way to export to them, because they're not directly supported in Photoshop. APNG and MNG are formats related to PNG, and support multiple still images as frames, in the same way as GIF, but don't have many of its limitations. Gimp can export to MNG natively, and APNG with a plugin. Then there's Google's WebP format, which aims to provide smaller files than PNG, and is an open standard—there's a Photoshop plugin. The problem with these new formats is that they're not universally supported, although recent versions of major browsers shouldn't have a problem.

Finally, things such as HTML5 Video or the dreaded Flash can display animations, and can be coded directly into websites. These have better support, but exporting from Photoshop can be trickier.

BUILD IT

ZAK STOREY, DEPUTY EDITOR

Full HD AMD Gaming Rig

So, this is what happens when Team Red accidentally sends you two RX 590s…

LENGTH OF TIME: **1–3 HOURS**

LEVEL OF DIFFICULTY: **EASY**

THE CONCEPT

WE'LL BE HONEST, even if AMD hadn't sent us two cards to test "by accident," we would have probably still had a crack at stacking an RX 590 into our monthly "Build It." It's an interesting release, at least on paper, if only because it doesn't make a whole lot of sense. We gave the RX 580 a lot of flak when it first launched, due to the fact that, spec for spec, it was identical to the RX 480, albeit with a little extra clock speed. Fast-forward another year, and bam! We're witness to a card with the same internal hardware specifications (although, apparently, it's a 12nm card now), and a slightly higher clock speed. Yay! Just kidding—it sucks. This isn't progress, but more on that later....

So, our hatred of obvious rebadges aside, AMD is the name of the game with this build. We do enjoy a bit of the ol' Ryzen love, and as we have a swanky new ROG Strix X470-F Gaming in house, along with a few more budget components, stacking them together seemed like an obvious choice. It's time to produce something positively mid-range, and hopefully 1080p-busting, too. Complete with a 250GB Samsung 860 Evo, a 1TB WD Blue, and 16GB of DDR4, this hex-core system should kick some serious booty in the world of the 1080p gamer.

COOL WHIP

LET'S BREAK IT DOWN into those sweet components. The big one is the Ryzen 5 2600X. This no-nonsense chip from AMD packs six cores, 12 threads, plenty of cache, a base clock of 3.6GHz, and a boost clock (single-core only) of 4.2GHz. It also comes with a cooler, but we're going with the not-so-low-key Cooler Master MasterAir G100M RGB, because it's relatively affordable, RGB, and looks darn baller.

On top of that, we're packing in two sticks of 8GB G.Skill SniperX 3,200MT/s in urban snow camouflage, so there's no chance anyone will see it storing our volatile information, at least during the winter months in those built-up areas. We've also added a 250GB Samsung 860 Evo SSD and a 1TB Western Digital Blue, to back it up with something a little more concrete. It's the go-to gaming staple now days, but still not something to be sniffed at. 250GB of solid-state flash for the operating system will see our rig performing smoothly, while the additional 1TB of slower spinning platters will keep our game library running nicely. We would see it as a media library, too, but as more of us in the office are turning to streamed sources of both music and films, we do wonder just how accurate a description that is anymore.

The weak link in the chain has to be the case. The BitFenix Enso is a fairly budget ATX tower that has some nice frilly RGB features embedded into the front panel, and the usual plethora of modern-day creature comforts littering the interior, but it has one massive problem: airflow. In short, the entirety of the front panel is solid, from top to bottom. The only air intake at the front is a wedge-shaped cut-out, located below the feet, running up to the front. It's not an intuitive design at all, so you have to either be in a very cool climate, be touting some serious air conditioning, or just not care particularly about thermals.

INGREDIENTS

PART		STREET PRICE
Case	BitFenix Enso Mid Tower	$80
Motherboard	Asus ROG Strix X470-F Gaming	$185
CPU	AMD Ryzen 5 2600X	$219
CPU Cooler	Cooler Master MasterAir G100M RGB	$37
GPU	XFX Radeon RX 590	$280
Memory	16GB (2x 8GB) G.Skill Sniper X @ 3,200MT/s	$138
PSU	600W Be Quiet! Pure Power 11 CM Gold	$100
Storage 1	250GB Samsung 860 Evo 2.5-inch SSD	$56
Storage 2	1TB Western Digital Caviar Blue 3.5-inch HDD	$46
OS	Windows 10 Home 64-bit	$100
Total		**$1,241**

1

OUTSIDE THE BOX

WE DON'T ALWAYS do it ourselves, but if you're building a rig, it's good practice to test the main components outside the chassis first. Use your motherboard box as an antistatic test bench, carefully lay the mobo itself on top, and start installing your hardware. We've taken this opportunity to carefully install our memory, CPU, and cooler first, because it's easier to do this outside the case than in. The MasterAir's a funny old thing; the mounting solution, although secure, is difficult to access. There's a wrench included to turn the top nuts, but we found it far easier to tighten them by hand.

ROG STRIX
X470-F GAMING
SNIPER X

2

NO FERRARI

OUR NEXT STEP was to strip the Enso down as much as we could. This included yanking the rather stubborn front panel off, along with the side panels, before becoming embroiled in a battle with the "Aura Sync" sticker stuck to the side of the PSU shroud. It's all fun and games until it leaves a nasty white residue behind when you remove it, and you spend the next 30 minutes trying to clean it off with lighter fluid and a microfiber cloth. As an aside, we're also hauling that front 120mm fan out of the chassis, and replacing it with two 140mm fans we had lying around, in an attempt to bolster the internal airflow.

3

MOTHERBOARD IN

THANKFULLY, BECAUSE WE'VE already pre-installed the CPU cooler and RAM, installing the motherboard into the chassis is a relatively painless procedure—just make sure you find the stand-offs, which are hidden in a little clear bag along with the case screws, tied to one of the hard drive caddies, first. Fortunately, the ROG Strix X470-F has its backplate pre-attached to the motherboard, so there's no forgetting that.

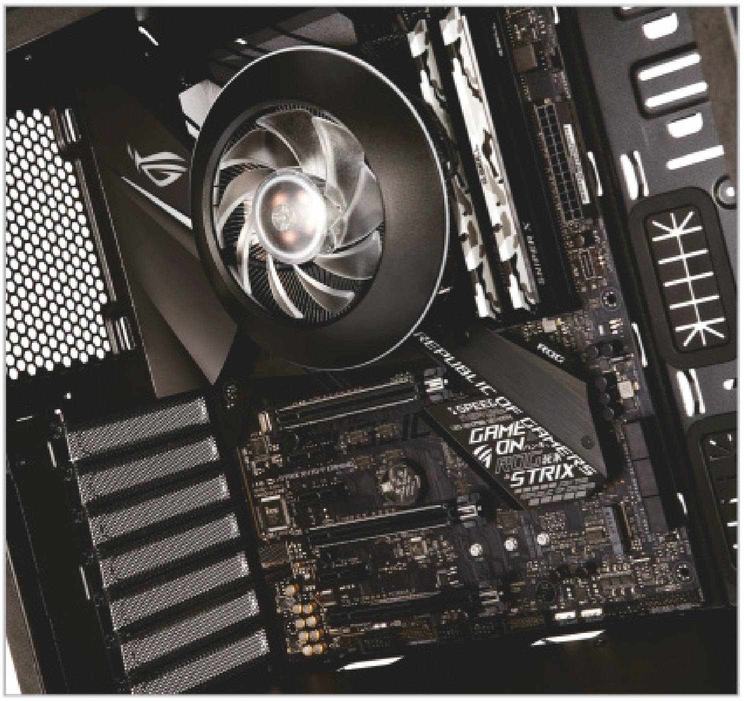

4

SHHHH!

NEXT UP ON THE LIST is the PSU. In this case, we're running with a 600W Be Quiet! Pure Power 11. This new PSU from the German company is a semi-modular beauty, coming in at an impressively affordable price point. Only downside? The EPS power is a bit short, and the case isn't that great with routing, so we've had to thread the power cable up and through the CPU socket cut-out. To do this, you have to remove the screws that secure the motherboard, run your cable through the cut-out, plug it in, then resecure the motherboard back down. It's a bit tight, but will run absolutely fine.

5

SOLID-STATE SNEAKINESS

WE'VE ALSO ATTACHED our SSD to the back of the motherboard tray, with the included caddy located here. There's only one as standard, but it's a fairly painless experience. It may sound obvious, but make sure you use flat-ended SATA cables for your 2.5-inch devices here. Also worth pointing out is that, although you can technically fit a 2.5-inch HDD or hybrid drive, it's not advisable, because the orientation can lead to additional stress on the drive and more overall noise, due to vibrations against the rear side panel. Instead, use the caddies located in front of the PSU, for 3.5-inch drives.

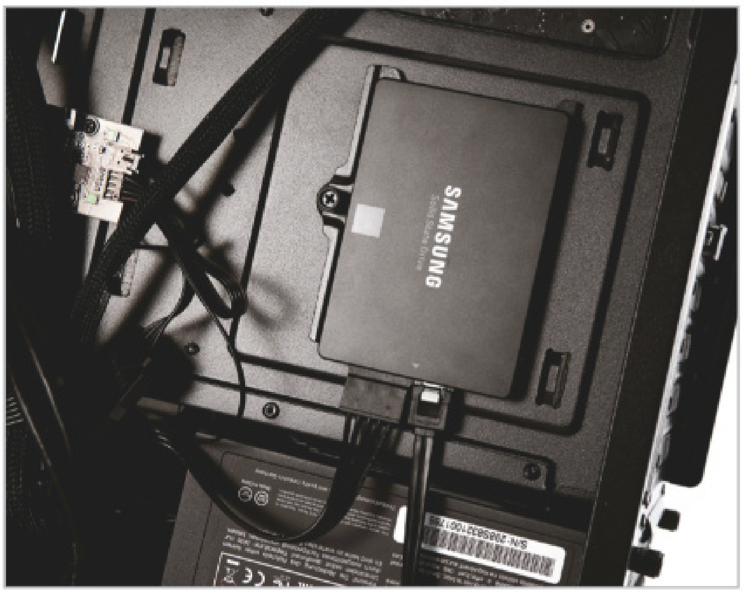

6

RADEON FOREVER

LAST BUT NOT LEAST is the centerpiece of our build: the XFX Radeon RX 590. It's a beast of a card, and looks absolutely stellar in our system. From the angled GPU block to the carbon fiber and black ventilated backplate, it certainly looks the part, and is really clean, too. We also took the opportunity to tie the two PCIe power cables together, before cutting the loose end off; because it's all black, it should look fairly inconspicuous, and keep everything tidy and together. The last thing to do is plug in the front I/O headers, and we're good to go.

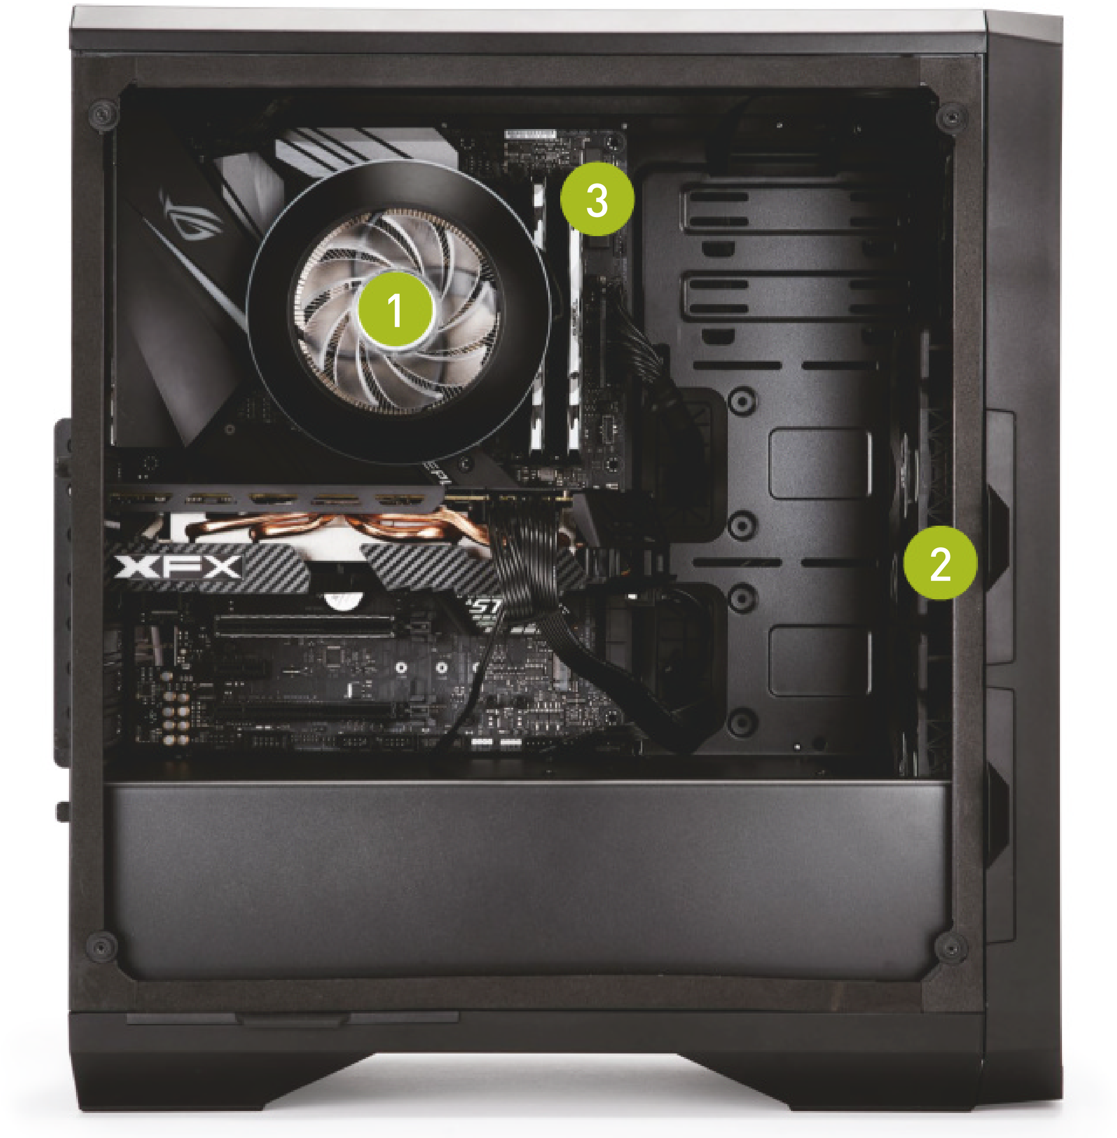

1 The MasterAir has two connections: one to sync the RGB LEDs with the motherboard, which we've connected at the bottom-left; and another, top-right, for the fan.

2 Three 120mm fans may have been a better solution for front intake, instead of two 140mm; unfortunately, the mounting holes dictate you can only install two 140mm fans, as we've done.

3 Crucial Ballistix memory, or perhaps something a little less camouflaged, may have been a better option for this one.

THAT'S A WRAP

RIGHT, THAT'S OUR RX 590 build done and dusted. A $1,250 budget 1080p build, with color-coordinated sleekness, tempered glass, and a minimal number of classy RGBs. Is it perfect? No. The case is the weak point, and although the design is fairly understated and classy in contrast to many at this price, the thermal limitation from the front panel leaves a lot to be desired.

Building inside the Enso was painless, though; the only trouble we had was with the sticker on the side of the panel, and trying to find the bag containing the screws and stand-offs. If we were upgrading, it would be one of the first things we'd change, perhaps alongside a better cooling solution.

Overall, performance is where you'd expect. The RX 590, coupled with the 2600X, makes this a prime-time 1080p gaming machine, with most AAA titles easily hitting the high 50s. The Samsung 860 Evo also performs admirably in file transfers, and the Ryzen 5 2600X hammers those Cinebench scores nicely, too, with 1,378 points in multicore, and 173 in single-core.

One of AMD's marketing pitches for the RX 590 was that over 64 percent of the PC gaming audience still plays games at 1080p. The RX 590, then, would be perfect for them, as it performs very strongly at this resolution. (And let's confirm here: It is not a 1440p card, but a 1080p card.) The problem with that is that, yeah, 64 percent may play games at 1080p, but they likely already have a card that's capable of driving games at 1080p. After all, it's not a new resolution, by any means, and there have been cards available for a number of years that are more than capable of driving all of today's titles at that resolution. So, why develop another rehash of a rehash to target a market that is already satisfied with the GPUs they have, and is not even mind-blowing from a price perspective?

The worst part is, it's not as though cryptocurrency is going to pick up the slack now either. With that market seemingly dead, and it being rumored that even Nvidia is going to need two years to shift the remainder of its GTX 1060 stock, well, you get the picture.

All in all, this rig is a fantastic entry point into the mid-range 1080p gaming ecosystem. And don't get us wrong—if you're moving up to that resolution, or haven't had the chance or the coin in the last few years to jump on to a modern GPU, the RX 590 is a good card to go for. But to design an entirely new product off 1080p just seems foolish.

BENCHMARKS

	ZERO-POINT	
Cinebench R15 Multi-Thread	1,152	1,378 (20%)
CrystalDisk QD32 Sequential Read (MB/s)	3,400	542 (-84%)
CrystalDisk QD32 Sequential Write (MB/s)	1,720	514 (-70%)
Rise of the Tomb Raider (fps)	60	64 (7%)
Total War: Warhammer II (fps)	46	49 (7%)
Tom Clancy's Ghost Recon Wildlands (fps)	39	43 (10%)
3DMark: Fire Strike (Index)	11,101	11,994 (8%)

0% 10% 20% 30% 40% 50% 60% 70% 80% 90% 100%

Our zero-point consists of an AMD Ryzen 5 1600, 16GB of Crucial Ballistix Sport LT @ 2666, an EVGA GeForce GTX 1060 3GB, and a 250GB Samsung 960 Evo M.2 PCIe SSD. All tests were performed at 1440p at the highest graphical profile.

CRAFT, BUILD, FIGHT, AND EXPLORE IN THIS BRAND NEW MINECRAFT SPECIAL!

Packed full of build guides, gameplay advice, and tips from your favourite YouTubers,
The Ultimate Guide to Minecraft is an essential companion for any Minecraft fan

132 PAGES OF BRAND-NEW MINECRAFT FEATURES
VOL 29 SEPTEMBER 2018
GAMES MASTER PRESENTS
100% UNOFFICIAL!
THE ULTIMATE GUIDE TO MINECRAFT
PS4 • XBOX ONE • SWITCH • PC • PS
S • ANDROID
SURVIVE A DAY IN ANCIENT EGYPT
Experience 24 hours in the land of the Pharaohs, among mummies and pyramids
+
10 SAND AND SNOW MAPS
REDSTONE SECRET STAIRCASE
ESSENTIAL TIPS FOR DESERT LIFE
10 DESE STRUCTU
FROM SNOW CITI TEMPLES, DISCO SOME AWESOM BUILDS
SURVIVE THE DESERT
THE ULTIMATE SAND & SNOW ADVENTURES!
FEATURING
ROMAN GARDEN
Create super-sized architecture
FANTASY ICE CASTLE
Construct a giant fortress
ABANDONED ARENA
Build columns, seating and more
FUTURE
VOL 29
ON SALE NOW

FUTURE

Ordering is easy. Go online at:

www.myfavouritemagazines.com

Or get it from selected supermarkets, newsagents, or bookstores

REVIEWS

TESTED. REVIEWED. VERDICTIZED.

INSIDE

74 Sapphire Radeon Nitro+ RX 590 8GB

76 AMD Threadripper 2920X

79 AMD Athlon 200GE

80 Asus ROG Strix GL12CX COD Edition

82 Acer Predator X27

85 Crucial BX500 120GB

86 Audio-Technica ATH-M50xBT

88 Fnatic Clutch 2

89 Thermaltake Tt Premium X1 RGB

90 Call of Cthulhu

91 Hitman 2

92 Lab Notes

AMD
RYZEN
THREADRIPPER

AMD THREADRIPPER 2920X PAGE 76

SAPPHIRE
NITRO
NITRO
HDMI
DVI
HDMI
SAPPHIRE

Wow, is that thing blue!

Sapphire Radeon Nitro+ RX 590 8GB

Look who's back, back again

YOU KNOW WHAT we don't like doing? We don't like reviewing the same product over and over and over again. The RX 480 was a fantastic breath of fresh air from AMD when it arrived. It showed promise. A mid-range GPU with enough clout to take on the 1080p battleground, no sweat, and at an affordable price, too. OK, it may have drawn perhaps a touch too much power from the PCIe bus at first, but that had little effect on premium motherboards anyway. It was an exceptionally strong GPU.

Fast-forward a year, and we saw AMD's first ever like-for-like rebadge, the RX 580. With Vega not having dropped yet either, the future didn't look good. And, lo and behold, it wasn't. Usually, a brand new top-tier architecture, like Vega, would populate the top two cards of a series, with the preceding architectures shimmying down the chain by one numerical designation each to fill in the gaps. A 380 would become a 470, a 480 would become a 570, you get the picture. In this case, though, nope. The RX 480 became the RX 580, with slightly higher clocks and power draw.

So, what's the RX 590? A new GPU? Nope. A new architecture? Nope. More transistors? Nope. Better performance? Marginally. Better price? Sort of. What's the point? Good question. We're not too sure. This is essentially a 12nm variant of the RX 480, a card that launched over two years ago, with 10 percent higher clocks. It's frustrating. Still, we gave AMD the benefit of the doubt, and listened to its marketing spiel. This time around, it's all about capitalizing on that 1080p market. Yup, according to AMD's research, 60 percent of gamers game at 1080p. That's a big portion. Because of that untapped potential, AMD decided in its wisdom that the time was right for an ever-so-slightly better GPU to enter the market, taking advantage of a 12nm Polaris architecture. You know, the same architecture that already targeted this market two years ago, but with a few extra frames.

MISSING MARKET

There's a problem with this. If 60 percent of the market is already gaming at 1080p, they probably already have a suitable setup. Especially as it's a resolution that's been comfortably accessible by mainstream gamers for well over a decade now. Hell, even if we ignore the massive market share Nvidia currently has, and just look at AMD's offerings over the last three years, the number of cards aimed at 1080p is a total of four—now five with the 590. Include Nvidia's GTX 1050 Ti, 1060, and upcoming 2060, and, well, you get the point.

That aside, the RX 590 does still perform well at 1080p. *Warhammer II* and *Ghost Recon* both scored in the mid-40s range, with *Shadow of War* and *Tomb Raider* raising that up to the 70–80 frames-per-second mark. It's a solid 1080p card. If nothing else, AMD has certainly achieved what it set out to do—a commendable effort.

Temperatures are well within acceptable parameters in our 21 C office, and the card is quiet overall. Sapphire has done a great job with the cooling on this card. The color scheme, however, leaves a lot to be desired. In a world of RGB and color-neutral cards, having a baby-blue GPU just doesn't fit with any of the systems out there today.

Where does all that leave us? With a slightly higher clocked card that performs slightly better than a "flagship" mid-range GPU that AMD launched over two years ago. It's not enough. We need a viable, competitive solution, to really take the fight to Nvidia, not just sit there wagging its tail and enjoying dominance in the world of consoles. It might as well be a monopoly otherwise. Come on, AMD, pick up the slack! **-ZAK STOREY**

VERDICT
7

Sapphire Radeon Nitro+ RX 590 8GB

+ BORING Good 1080p performance; mediocre pricing.

− POINTLESS No market for it; horrific color scheme.

$280, www.sapphiretech.com

BENCHMARKS

	Sapphire Radeon Nitro+ RX 590 8GB	Nvidia GeForce GTX 1060 6GB
Middle Earth: Shadow of War (Min/Avg fps)	**35/83**	30/53
Total War: Warhammer II (Min/Avg fps)	**40/47**	35/43
Ghost Recon Wildlands (Min/Avg fps)	**36/42**	**36**/41
Rise of the Tomb Raider (Min/Avg fps)	**16/72**	**16**/55
Unigine Superposition (Min/Avg fps)	**17/21**	14/17
3DMark Fire Strike (MB/s)	**14,241**	11,505
Power Draw Idle/Load (Watts)	107/352	**71/227**
Price to Performance (Index)	**$7.00**	$7.28

Best scores are in bold. Our test bed consists of an Intel Core i7-8700K, 16GB of G.Skill DDR4-3200, an MSI Z370 Gaming Pro Carbon, and a 500GB Samsung 960 Evo M.2. All games are tested at their highest graphical profile, at 1080p.

SPECIFICATIONS

GPU	Polaris
Lithography	12nm
Transistor Count	5.7 billion
Shader Cores	2,304
Texture Units	144
ROPs	32
Core/Boost Clock	1,469/1,545MHz
Memory Capacity & Type	8GB GDDR5
Memory Speed	8GT/s
Memory Bus	256-bit
TDP	225W
Display Connectors	DisplayPort, HDMI, DVI-I

AMD Threadripper 2920X

Twelve cores of pure unadulterated joy

WHEN THE FIRST SAMPLES of the early Threadripper processors ventured out into the wilds of the technology journo landscape, we were blown away. The high-end desktop processors were unlike anything we'd ever seen. They were incredible, with 16 cores and 32 threads, all for a little under a grand. AMD had slapped Intel's high-end CPUs in the face, and made them take a long hard look at themselves in the mirror. We'd expected as much. The Infinity Fabric design, the modular core complexes, the EPYC server arsenal—it was all lining up to create a perfect storm in which these high-end consumer parts could make some serious headway into the market.

The thing is, they weren't perfect. They had their niggles. Memory compatibility, latency issues, slow single-core clock speeds—the Threadripper lineup was hardly bereft of problems. If plug-and-play was the name of the game, Intel was still king, and the Core i9-7900X—although 12 threads down in comparison—still provided the best option for the high-end, power-hungry, plug-and-play enthusiast.

A year later, and like a good cheese, the Zen architecture has finally begun to mature. Threadripper 2, or "Plus," or whatever the marketing types want to call it, has arrived. With it, a round of optimizations has followed, addressing that pesky memory latency issue. Has it made a difference? Absolutely. Throughout our Threadripper 2 testing, we encountered minimal issues in that department. The BIOS bugs and memory latency problems that plagued the platform upon launch are now a thing of the past. To put that into perspective, our Threadripper 2920X scored a solid 62.5ns memory response time. Intel's Core i9-7980XE, on the other hand? 88ns. The Threadripper 1950X today? 87.6ns. When we tested it originally? Over 100ns. That's a staggering improvement over the course of a year, and with Zen 2 featuring an even better memory design, it's definitely going to be one to watch.

LIVING THE DREAM

Memory aside, all of this pushes the 2920X closer toward being that plug-and-play part we dreamed of back in 2017. Its 12 cores and 24 threads provide a fantastic base for streaming, rendering, and gaming. You don't even have to worry about setting up Ryzen's "game mode" to ensure your memory is operating correctly in game, either. On top of that, temps are incredible. Under our 240mm full-cover AIO, our review sample topped out at 60 C under load, and even after bumping up all 12 cores to 4.3GHz (1.42V), the most we saw was 72 C. That's simply insane.

The competition, ironically, comes from AMD's own lineup. In particular, the previous-generation Threadripper 1950X. The 2920X costs $650; you can now pick up a 1950X for $55 less. That's four more cores for less than half the price of this reviewer's weekly groceries; Zak loves yogurt (and if you don't get that reference, he's leaving the mag). You do lose out on that bumped-up single-core performance, and those sweet memory optimizations, but when it comes to multi-threaded rendering, you're in for a treat. In reality, you have two processors aimed at two separate audiences, with two different specifications, which, although they kinda compete with each other, they also kinda don't.

When AMD set out to bring Threadripper 2 to market, its PRs told us it had two different parts: the WX, for workstation professionals, designed as a no-holds-barred memory rendering monster; and the X, for gamers, enthusiasts, and streamers. At first, we didn't believe them. After all, despite its niggles, the first gen was fine for both, and marketing speak is usually just that: marketing. Now? Well, it's nothing if not impressive to continually see AMD deliver on its CPU promises. **–ZAK STOREY**

PROCESSOR BENCHMARKS

	AMD Threadripper 2920X	AMD Threadripper 1950X
Cinebench R15 Single (Index)	**176**	166
Cinebench R15 Multi (Index)	2,575	**3,010**
Tech ARP's X264 (fps)	54.14	**64.57**
Fry Render (m:s)	1:39	**1:23**
AIDA64 Memory Latency (ns)	**62.5**	87.6
AIDA64 Memory Bandwidth (MB/s)	**75,985**	75,412
Power Draw Idle/Load (Watts)	**69**/272	91/**271**
Total War: Warhammer II (fps)	**75**	**75**
Ghost Recon Wildlands (fps)	**58**	56
Middle Earth: Shadow of War (fps)	**94**	84
3DMark Fire Strike (Index)	16,909	**18,899**
Maximum Overclock @ Voltage (GHz @ V)	**4.3 @ 1.42**	4.0 @ 1.4
Price to Performance (Index)	3.96	**5.06**

Best scores are in bold. Our test bed consists of an Asus ROG Zenith Extreme, 4x 8GB of Corsair Dominator Platinum, and an Nvidia GeForce GTX 1080. All games were tested at 1080p on the highest graphical profile, with average fps noted.

VERDICT
9
KICK ASS!

AMD Threadripper 2920X

+ ROPE BURN Impressive all-around performance; very strong pricing; incredible thermals; huge jump in memory over previous gen.

− TANGLED TWINE Outpriced by the 1950X.

$650, www.amd.com

SPECIFICATIONS

Base/Turbo Clock	3.5/4.3GHz
Cores/Threads	12/24
Lithography	12nm
Cache	L3 32MB
Memory Support	128GB DDR4
Memory Channels	4
Max PCIe Lanes	64
TDP	180W

AMD Ryzen Threadripper 2920X
AMD
RYZEN
THREADRIPPER
YD292XA8UC9AF
UA 1840SUT
9HG9140V80124
DIFFUSED IN USA
MADE IN CHINA
© 2017 AMD

You know it's bad
when you're your
own competition.

MAKE YOUR OWN PROJECTS WITH THE RASPBERRY PI

Learn the electronics, computing, and coding skills you need to make your own projects with the Raspberry Pi, and let your imagination run wild

Practical
NEW
Raspberry Pi Projects
Get hands-on with your Raspberry Pi
Assemble a robot arm
Create a virtual reality setup
Drive a remote control car
Over 60 projects
Program a RasPi-powered synth
Control Minecraft with a power glove
Get to know the world's smallest computer
Fly your own quadcopter
Attach screens and other accessories
CUSTOMISE YOUR PI
BUILD INCREDIBLE GADGETS
PROGRAM AMAZING SOFTWARE
ISBN 978-1-78389-741-4
FUTURE
FOURTH EDITION
Practical Raspberry Pi Projects Fourth Edition

ON SALE NOW

FUTURE

Ordering is easy. Go online at:

www.myfavouritemagazines.com

Or get it from selected supermarkets, newsagents or bookstores

AMD Athlon 200GE

The budget chip you've been waiting for? That depends…

CONVENTIONAL WISDOM says that AMD is the budget darling, while Intel has you covered at the high end. Conventional wisdom is out of step with reality. Because, right now, if you want serious performance, AMD's Threadripper should be at the top of your list, while anyone on a super-tight budget is better served by Intel's Pentium chips and the likes of the Core i3-8100. AMD doesn't want to leave any market segment to Intel, though, so has revived the Athlon brand for this, its most affordable Zen chip to date.

Before we delve into the details, it's worth mentioning the price. In theory, this is a $55 CPU, but in truth, you'll do well to actually find it for that, because it's one of those chips that tends to sell out. By way of example, the cheapest we could find it for on Newegg was $67, which is still a great price, if not quite the ridiculous bargain that it should be. Note that this paltry sum also includes a CPU cooler, and in our testing, it is indeed every bit as "whisper quiet" as AMD claims. So, either way, you get a lot of hardware for your money.

Apart from a rather unimpressive-looking, yet still capable, CPU cooler, what does your small chunk of cash get you? This is a dual-core, quad-thread processor, with a 3.2GHz base clock, and Radeon Vega 3 integrated graphics. It uses the same AM4 socket as the bigger Ryzen chips, has a low, low TDP of just 35W, and supports dual-channel DDR4 memory. The integrated graphics aren't going to have you playing the latest games at maximum settings, but with sufficient tweaking, you should get some classics and efficient modern titles running at passable frame rates.

We've actually already featured the Athlon 200GE in *Maximum PC,* as a few issues ago, we used it as the basis of our $285 "Build It" challenge. At the time, we were blown away by its obvious value, but as that machine didn't boast a discrete graphics card, it wasn't the whole story in terms of true CPU performance. Now it's in our test bench, we can compare it with the field, and see where it excels and where it pulls up a bit short.

ON THE TABLE

A quick look at the benchmark table below reveals that it performs pretty much where we'd expect it in most CPU tests. Where it does falter is when it's paired up with a decent graphics card, in the form of the Nvidia GeForce GTX 1080, and this is because the integrated graphics butcher the PCIe connectivity down to just 8x. The 3DMark score is significantly lower than non-APU chips, and the gaming performance is also off the pace. If you're looking to get into gaming on the cheap, this is essentially a false economy, and you'd be better off going for a Ryzen or Core i3.

On the plus side, this is a low-power chip, and even under load, it doesn't draw much from the wall. It also offers great value for money—even at $67, it boasts better bang for your buck than most chips. If you're looking for a small system that is going to make the most of the integrated graphics, for a media center, or a discreet server, then there's a lot to like here. Gamers, serious users, and anyone looking for an affordable base system for bigger and better things are going to be ultimately disappointed, unless you factor in a CPU upgrade as part of it—it does use the standard AM4 socket, after all. **–ALAN DEXTER**

BENCHMARKS

	AMD Athlon 200GE	Intel Core i3-8100
Cinebench R15 Single (Index)	121	**153**
Cinebench R15 Multi (Index)	348	**598**
Tech ARP's X264 (Avg fps)	7.81	**14.76**
Fry Render (m:s)	11:57	**7:28**
AIDA64 Latency (ns)	84.1	**63**
AIDA64 Memory Bandwidth (MB/s)	37,859	**42,076**
Power Draw Idle/Load (W)	44/**67**	**43**/112
Total War: Warhammer II @ 1080p (fps)	57	**82**
Ghost Recon Wildlands (fps)	36	**63**
Middle Earth: Shadow of War (fps)	75	**96**
3DMark Fire Strike (Index)	10,207	**15,237**
Price to Performance (Index)	**5.19**	4.6

Best scores are in bold. Our test bed consists of an Asus Crosshair VI Hero, 2x 8GB of Corsair Dominator Platinum, and an Nvidia GeForce GTX 1080. All games were tested at 1080p on the highest graphical profile, with average fps noted.

VERDICT

8

AMD Athlon 200GE

ZEN Great value for money; low power; versatile; reasonable integrated graphics.

REBELLION Hobbled PCIe lanes; not overclockable.

$67, www.amd.com

SPECIFICATIONS

Base/Turbo Clock	3.2GHz
Cores/Threads	2/4
Lithography	14nm
Cache	L3 4MB
Memory Support	64GB DDR4
Memory Channels	2
Max PCIe Lanes	8
TDP	35W

Asus ROG Strix GL12CX COD Edition

A super PC for Call of Duty super-fans? Hmm…

THIS ASUS ROG Strix GL12CX combines the dynamism of the RTX 2080 and the Intel i9-9900K. What's special about this power couple is that the CPU comes factory overclocked up to 4.9GHz under a dead silent, liquid-cooled system, which seems like a good investment if you're the kind of person who is into high-end tech.

Turned off, the solid black chassis captures the silent ninja look Asus was presumably going for. Turned on, small strips of color glow brightly from between the cracks. This special *Black Ops 4* edition also comes with a clear side panel to show off the components. Aura Sync lighting adds a specific color pattern to your keyboard as soon as you start *Black Ops 4*, then highlights specific keybindings once in game. It's pretty, but not exactly essential. One of the wonkier features is the magnetized front panel; it provides easy access to the optical drive and hot-swappable SSD tray, but doesn't attach to the case. The whole thing pops off, which is a little awkward.

At 1080p on Ultra, the RTX 2080 handles some of the most demanding games with ease. Whether we were struggling to survive against an onslaught of zombies or navigating our way through a busy crowd in Mexico, not once did our screen suffer any kind of frame rate dip. The in-game benchmarks showed *Shadow of the Tomb Raider* and *Total War: Warhammer II* both averaged consistently between 90 and 100fps. *Tom Clancy's Ghost Recon: Wildlands* came in at an average of 65fps; lower, but still within an ideal range. It did push the GPU to its limits, though—up to 97 percent capacity—which brings us nicely to the 3DMark benchmarks.

KNOW THE SCORE

We didn't have a 1440p or 4K monitor to hand to run the gaming benchmarks at higher resolutions, but we did test the GPU in Fire Strike Extreme and Ultra. The RTX 2080 scored 11,683 for Extreme (1440p) and 6,159 Ultra (4K), which is about in line with—perhaps slightly lower than—what you should expect from Nvidia's new GPU.

The i9-9900K provides great support to the RTX 2080. Looking at the single-core benchmarks in Cinebench, this overclocked CPU comes in at an average score of 212, which is higher than the top CPU on its list. However, the multicore test only brought the i9-9900K in at an average of 2,090, which is lower than the i9-7900X. Obviously, take this with a grain of salt; 3DMark puts the CPU usage at about 20 percent on the low end and 52 percent on the high end for *Ghost Recon: Wildlands*, so the i9-9900K can handle way more in terms of gaming.

This GL12CX has an MSRP of $3,799; if you look at the current retail price of each piece of hardware individually, you could save roughly $500–$1,000 by building an equivalent system yourself, but that's without factoring in the cost of prebuilt benefits such as warranty, labor, and, in this case, the special *Call of Duty* edition extras. There's also a cheaper version with a factory overclocked i7-9700K, RTX 2080, and 16GB RAM for $2,600.

Speaking of extras, those only include a free copy of the game and a bunch of stickers. It's also questionable whether a keyboard and mouse are included. The Asus store lists these as add-ons, but while some retailers say one version of this desktop comes with a keyboard and mouse, others say they are sold separately. The peripherals don't have a special *COD* spin on them either, but if you're in need of a new mouse and keyboard, they're a nice addition.

The GL12CX is nice, but overpriced (like most prebuilts), especially if you don't care about the *COD* extras. If you're in the market for a quiet and powerful prebuilt, it might be worth considering—just don't make the stickers and free game your deciding factor. **–JOANNA NELIUS**

VERDICT 7

Asus ROG Strix GL12CX COD Edition

+ TO THE DEATH Compact, near-silent design; hot-swappable SSD tray.

− CALL OF DUTTIE Gimmicky stickers; price.

$3,799, www.asus.com

BENCHMARKS

	ZERO-POINT	
Cinebench R15 (Index)	1,152	2,090 (81%)
CrystalDisk QD32 Sequential Read (MB/s)	3,400	2,947 (-13%)
CrystalDisk QD32 Sequential Write (MB/s)	1,720	1,664 (-3%)
3DMark: Fire Strike (Index)	11,101	22,111 (99%)
Rise of the Tomb Raider (fps)	37	91 (146%)
Total War: Warhammer II (fps)	46	103 (124%)
Tom Clancy's Ghost Recon: Wildlands (fps)	39	66 (69%)

0% 10% 20% 30% 40% 50% 60% 70% 80% 90% 100%

Our zero-point consists of an AMD Ryzen 5 1600, 16GB of Crucial Ballistix Sport LT @ 2,666, an EVGA GeForce GTX 1060 3GB, and a 250GB Samsung 960 Evo M.2 PCIe SSD. All tests were performed at 1080p at the highest graphical profile.

SPECIFICATIONS

Processor	Intel Core i9-9900K up to 5.1GHz 8-core
Graphics	Nvidia GeForce RTX 2080 8GB GDDR6 VRAM
RAM	32GB (2x 16GB) DDR4 SDRAM @ 2,66MHz
Motherboard	mATX Z390
Primary Storage	512GB NVMe M.2 SSD
Additional Storage	2TB 7,200rpm HDD
Cooling Solution	HyperStream liquid cooling
PSU	700W
Case	GL12CX
Warranty	One year

It looks better without the stickers.

Acer Predator X27

HDR done right on the PC, but at a price

Color Accuracy
NVIDIA
G-SYNC HDR

THE ARRIVAL OF HDR (high dynamic range) image rendering on the PC hasn't exactly been a slam-dunk success. Take support for HDR within the Windows OS itself. Let's be generous and call it patchy. Ironically, the very notion of what it means to be an HDR display has also been murky, thanks to numerous conflicting standards.

The result has been a bit of a blowout. Can you even be sure what you're seeing is truly HDR? Often, it's been far from obvious. But not with Acer's awesome new Predator X27. When this new 27-inch monitor has its afterburners lit up in HDR mode, there's no doubt you're looking at something new, something special. This is HDR done right on the PC.

In fact, the X27 does a lot of things right. For starters, it's based on a quality 27-inch IPS panel that supports full 4K. To that it adds 144Hz refresh rate and several other frills, including Nvidia's G-Sync tech in HDR spec. But the showbiz feature is surely the X27's backlight. It's split into 384 dynamic zones. Yup, this monitor offers proper local dimming.

Combine that with a peak local brightness of 1,000 nits, and you have the makings of a monitor capable of true dynamic range rendering. Running the same content side by side with an SDR panel, even a high-quality 4K IPS model, the difference is little short of epic. There's so much visual fizz and pop on the X27. What's more, there's little to no glow around bright objects on dark backgrounds. Suddenly, even a top-notch SDR screen looks flat, dull, and dead.

Of course, whether it's video or games, genuine HDR content is relatively hard to come by. So, the X27's performance with SDR content is critical. To cut a long story short, you have two options. You could run the X27 as a straight SDR display, in which mode it's simply a very good conventional monitor, albeit one that offers the unusual attraction of combining a 4K native resolution with 144Hz refresh.

Enable HDR in Windows, however, and the X27 picks up the new signal format and switches modes. You can then use Windows 10's newly updated HDR settings to boost the vibrancy of what would otherwise be very dull SDR content, including the Windows interface. The outcome is a monitor that looks great with SDR and HDR content at the same time. That's important, as it means you don't need to constantly toggle between HDR and SDR modes in Windows.

Admittedly, a few glitches remain in Windows HDR support, but recent updates have improved it dramatically, and the X27 shows that progress off in style. That doesn't necessarily make this monitor easy to recommend, however. For starters, there's the price. Yes, this is a fantastic display, but it's one heck of a lot of money to pay for a mere 27-inch monitor. If you want something to watch HDR movies on, there's far more value to be had from the HDR HDTV market.

The X27's proposition as a gaming monitor isn't a no-brainer, either. We tested the X27 with Nvidia's latest uber-GPU, the GeForce RTX 2080 Ti. Even with that card, you're not going to see high enough frame rates to fully take advantage of the X27's 144Hz refresh in a lot of modern games. That isn't the X27's fault, but it begs the question of what, exactly, you're paying for. In the end, the X27 lives and dies on pricing. It's a truly fantastic display, but just like Asus's almost identical ROG Swift PG27UQ monitor, the X27 is simply too expensive to be anything other than a niche player in the HDR monitor market. –**JEREMY LAIRD**

VERDICT
9

Acer Predator X27

+ OPEN RANGE Stunning HDR visuals; awesome 4K resolution and 144Hz combo.

− RANGE ANXIETY Megabucks for 27 inches; 4K plus 144Hz equals massive GPU load.

$1,899, www.acer.com

SPECIFICATIONS

Panel Size	27-inch
Native Resolution	3840x2160
Pixel Density	163.18
Panel Type	IPS
Maximum Refresh	144Hz
Response	4ms G2G
Contrast	1,000:1 (panel native)
Display Inputs	DisplayPort, HDMI
Connectivity	USB
VESA Mount	100 x 100mm
Warranty	Two years

The X27's show-stopper is a 384-zone dynamic backlight.

XBOX
THE OFFICIAL MAGAZINE

ISSUE 219 NOVEMBER 2018
WWW.GAMESRADAR.COM/OXM
XBOX
OFFICIAL MAGAZINE

ASSASSIN'S CREED ODYSSEY
MASSIVE PREVIEW OF THE EPIC GREEK RPG

WE'VE PLAYED IT!
RESIDENT EVIL 2
First look at the stunning Xbox One remake of the survival horror classic

MAKE AN INDIE GAME
How you can become a game developer

47 NEW GAMES
DARKSIDERS III OVERCOOKED 2 MUTANT YEAR ZERO INDIVISIBLE NO MAN'S SKY TWIN MIRROR EARTHFALL VERMINTIDE 2 NIDHOG 2 TRIALS RISING SEKIRO AND MORE...

The essential magazine for all things Xbox

NEW ISSUE ON SALE NOW

ORDER DIRECT FROM MYFAVOURITEMAGAZINES.COM/GAMING

Crucial BX500 120GB

Crucial blows it with this bargain-basement SSD

WE'RE BIG FANS of Crucial's SATA SSDs. In fact, it's been our go-to brand in this segment of the solid-state market for years. Its MX500 is our default recommendation for those wanting a big, reliable, and reasonably fast SATA SSD.

So, when we say it's blown it with the BX500, we take no pleasure in bashing what is generally a very good brand. But hang on, what exactly is the BX500? It's a new budget 2.5-inch SATA SSD, and first impressions are promising.

For starters, it's based on the same 64-layer 3D TLC NAND memory as the MX500. In other words, it's not a QLC drive, with all the concerns over performance and durability that would imply. What's more, Crucial rates even this bottom-feeding 120GB drive at 540MB/s for reads and 500MB/s for writes. That's within 10MB/s of even the beefiest 2TB MX500. Impressive.

The BX500 also uses a variant of the SMI SM2258 controller—again, similar to the MX500. So, what's the problem? Our suspicions were first aroused by Crucial's failure to quote random rather than sequential performance numbers. Even the detailed data sheet omits random performance. Odd. The 40TB of claimed write endurance isn't exactly confidence inspiring, either.

A closer look at the SMI SM2258 chip reveals something important: It's the SM2258XT, the suffix indicating this is the so-called DRAM-less version. Put simply, omitting a DRAM cache is an easy way to save a few bucks. But at what cost to the drive's performance?

In terms of sequential throughput in synthetic benchmarks, the BX500 isn't too terrible. It hits the claimed spec for reads, while writes of 238MB/s are down on the claims, but still look tolerable. The synthetic 4K random performance doesn't look all that bad, either.

CLIFF NOTES

All of which demonstrates how synthetic tests don't always capture a drive's true character. During our pre-test routine, which includes fully filling the drive, then deleting the data, the true horror of what Crucial has done emerged. Copy a multi-gigabyte file to the BX500, and it looks dandy at first, humming along just below 500MB/s. But after about 1.5GB's worth of transfer, the rate falls off a cliff, then oscillates between 10MB/s and 80MB/s.

The initial drop almost certainly reflects the point at which the drive's SLC cache is exhausted. Inevitably on such a small, cheap drive, that cache is limited. Even then, the post-cache performance is hideous, a fact that probably reflects the further lack of a DRAM cache.

It's not clear how larger drives from the BX500 portfolio perform; it's also available as 240GB and 480GB. Likely, you get a bit more SLC cache with the larger drives that will obscure the poor performance a bit longer. But you'll hit the underlying crapiness soon enough.

This is for buyers who choose based on a combination of price and superficial specs. It's dirt cheap, the headline read and write speeds look decent, so why not save a few bucks? The impact of that saving is catastrophic. You can't buy a 120GB MX500, but the 240GB BX500 is just $10 cheaper than the 250GB MX500, the latter being a bit bigger, and a far more consistent performer.

We could argue you get what you pay for, so buyers only have themselves to blame. But it's reasonable to make assumptions about the basic quality of a drive from Crucial without needing to dig deep into the technical documents or even read a review. The BX500 falls short of those expectations. This drive is unworthy of the brand. **–JEREMY LAIRD**

VERDICT

5

Crucial BX500 120GB

+ BY GRABTHAR'S HAMMER! Very cheap drive from a proper brand; mostly good specifications.

– WHAT A POINTLESS SAVING! Minimal SLC cache and no DRAM kills performance.

$29, www.crucial.com

BENCHMARKS

	Crucial BX500 120GB	Samsung 860 Evo 500GB	Samsung 970 Pro 512GB
30GB Internal Copy (Seconds)	224	89	**29**
CrystalDisk 6 Sequential Read/Write (MB/s)	561/238	563/533	**2,788/2,331**
Anvil 4K QD 1 Read/Write (MB/s)	30.8/112	**49.8**/119	47.7/**141**
Anvil 4K QD 2 Read/Write (MB/s)	53/197	**98**/225	91/**281**
Anvil 4K QD 4 Read/Write (MB/s)	87/207	**182**/314	152/**524**

Best scores are in bold. Our test bench consists of an Intel Core i7-8700K, an Asus Maximus X Hero, an Nvidia GeForce GTX 1080, 16GB of Corsair Dominator Platinum DDR4, and a Samsung 850 Evo 250GB OS SSD.

SPECIFICATIONS

Capacity	120GB
Interface	SATA III
Control Protocol	AHCI
Controller	SMI SM2258XT
NAND Type	64L 3D TLC
Sequential Read/Write	540/500MB/s
Read/Write IOPS	N/A
Warranty	Three years

There's even a dimple on the left earcup so that you can orient them in the dark.

Audio-Technica ATH-M50xBT

Fine sounds with full freedom

THERE ARE FUNCTIONAL headphones, and then there are experience headphones. Let's not skirt around the subject for any longer than necessary: With the ATH-M50xBT, Audio-Technica delivers exactly the kind of sound stage you'd expect a company like that to deliver, and it does it hard. This is a set of brash closed-back cans that well and truly thumps through music, seemingly throwing impact straight down through your spine when you push the decibels up, literally throbbing along with the bass end. But where some bass-heavy cans allow that ketchupy low end to smother the high-end fries, A-T knows there's a balance to be found, and the company has dipped the M50 drivers in just the right amount of sauce. Highs fizz, and mids stand out if you go looking for them, even when the driving frequencies start to go low.

This balance isn't going to suit all ears. If you prefer not be battered by your bass, you're not going to be happy with the mix, no matter how high the levels of clarity may be. But if you're looking for headphones that would really rather you stopped whatever else you were doing and listened to them instead, because they are well and truly in charge, then your weird fetish is fully catered for. Clamp these on your skull and the isolation takes over—that's partially a result of the closed-back design, partially a result of those drivers, and partially (we suspect primarily) down to the serious amount of torque the stiff headband applies to each earcup.

Normally, we would spit and curse about tightness, but not here. In other headphones, we'd whine about the inadequate padding. Again, no. While it's true that Audio-Technica's attention to cushioning has not been quite as laser-focused as its work on the M50's drivers, these headphones can sustain long wear times without crushing. Again, it's a case of balance—soft materials, padding that's just shallow enough to introduce your ears to the speakers, but not quite tight enough to crush at least this reviewer's cartilage, and the right kind of hinge to sit everything nice and snug, without arguing with your skull about what shape it should be. All of this in a package that, should you wish to take full advantage of its Bluetooth portability, folds up beautifully, with a notched hinge that stops the arms flopping about until you're ready to release them.

About that Bluetooth. It adds a good bundle of bucks on to the price of the wires-only M50, and it's sometimes a little finicky about pairing with certain devices, but it does transform these cans from desktop dailies to proper street warriors. Design-wise, they're clean and thin, less ostentatious (and more listenable) than the likes of Beats, and the battery goes on and on. Controls are located on the left earcup, the included cable also includes inline controls for phones, and there's a mic built in if you're forced to take a call on the move. The only real qualm we have with the Bluetooth end of things, beyond those pairing niggles, is codec support—they run Bluetooth 5.0, and can work with AAC and aptX, but not aptX-HD, a bit of a shame for a set of cans that brings otherwise pro performance and design down to the consumer level.

Even though we've been forced to say at least one bad thing, the superlatives absolutely have the most weight when it comes to the Audio-Technica ATH-M50xBT headphones. This is a harmonious combination of design and engineering that sounds just great. Yes, if you're a purist who'll never listen through Bluetooth, you'll probably be looking for something flatter than the rather overexcited sound of these cans, and that's fine. They're not for you. But they are for us, even at the price. **–ALEX COX**

VERDICT

9

Audio-Technica ATH-M50xBT

⊞ **WIRELESS** All-around sound; great battery; solid engineering, brilliantly portable.

⊟ **TIRED** Some Bluetooth awkwardness; no aptX-HD.

$199, www.audio-technica.com

SPECIFICATIONS

Driver Type	45mm rare earth
Impedance	38 ohms
Frequency Response	15Hz–28KHz
Design Style	Closed-back dynamic
Microphone type	Condenser
Connectivity	Bluetooth 5.0, 3.5mm jack
Weight	11oz
Cord Length	4 feet

Fnatic Clutch 2

Esports or not, this is definitely, absolutely, a mouse

THERE'S SOMETHING about Fnatic's ethos (and our own frequently facetious nature) that really makes us want to make fun of the esports company at every possible opportunity. Indeed, we've done that before, much to the chagrin of our coffee-addled legal eagles. But not today. Today we look at the Clutch 2, and we're going to treat it dead seriously. Just you watch.

The human race has been blessed, in the majority of cases, with two hands. This is a mouse that, despite its otherwise reasonably balanced sculpting, is suited for the rightmost of those; a pair of thumb buttons has been placed on its left-hand side. These buttons are, at least in the context of mice, rather large and easy to press. Wider thumbs may press them more easily, perhaps even by accident; their action is rather soft and indistinct, but much of that action happens long after the click of their microswitches, so they are suitably quick to activate.

As befits a PC mouse, the Clutch 2 sports more buttons than just those. There is a mode-select button nestled in the upper ridge, deep-action and stiff to press, as such things should be. In front of that is a clickable scroll wheel, calibrated to depress with an actuation force that is balanced enough to prevent accidental central clicks, and engineered so that every stiff rolling action of one's middle finger can only be deliberate. Astride the wheel sit two face buttons, their microswitches (manufactured by Omron, and able to withstand 50 million clicks) positioned perfectly to allow them to be depressed, with very little force, anywhere from the front edge of the mouse to a point adjacent to that mode-select button. This will suit users with unusual hand shapes or interesting grips. There's little depth to these buttons; where their side-mounted pals activate quickly and continue, these activate with no travel and stop dead.

Both the left and right edges of the Clutch 2 are adorned with a stippled silicon coating in an attempt to coddle one's thumb and ring finger, and prevent slipping. The efficacy of this coating depends on the skin type of the user. We found it less than effective, certainly compared to the more texturized silicon coatings employed by other mouse manufacturers, but this will not be the case for those with a less relaxed style of holding. The top layer, we're told, is coated in oleophobic paint, a layer designed to repel hand sweat in times of mousing intensity, which is a nice touch.

There's a handful of configurable colored LED lights inside, for those whose hands need a little illumination, and a large area of slippery silicon beneath. As a mouse of very little weight—the Clutch 2 is identical in bulk as its slightly wider but otherwise identical counterpart, the Flick 2—this generous foot offers an extreme amount of maneuverability. Its sensor, a Pixart 3360, is regarded by many as one of the most accomplished optical sensors on the market, and only adds to that movement acumen. Through software, the sensor can be tweaked to increase its lift distance, or configure the resolutions assigned to the mode-select button, from sluggish slow to outrageously fast.

In summary, then, the Fnatic Clutch 2 is a computer mouse. It's unlikely to offend anyone who puts their hands on it, and it can suit just about any style of grip. It is, perhaps, a little more expensive than it should be, considering the market it is in—there are devices that gamers will find more exciting, and there are mice with prettier designs and additional buttons—but the simplicity and stiffness of the Clutch 2 make it perfect for those who want to concentrate on what they're doing without making mistakes. "It's a rough road to the top," says box-adorner "Rekkles." "You need to trust your equipment no matter what." Rekkles, as we can plainly see, is a sage. **–ALEX COX**

VERDICT 7

Fnatic Clutch 2

+ **SERIOUS** Straight-ahead design; well-weighted buttons; class-leading sensor.

− **SERIOUSLY?** Unexciting and expensive; right hand only.

$60, www.fnatic.com

SPECIFICATIONS

Sensor	Pixart
Sensitivity	12,000 cpi
Sensor Model	Pixart 3360
Polling Rate	1,000Hz
Programmable Buttons	Yes
LEDs	RGB
Cable Length	6ft
Weight	3.42oz

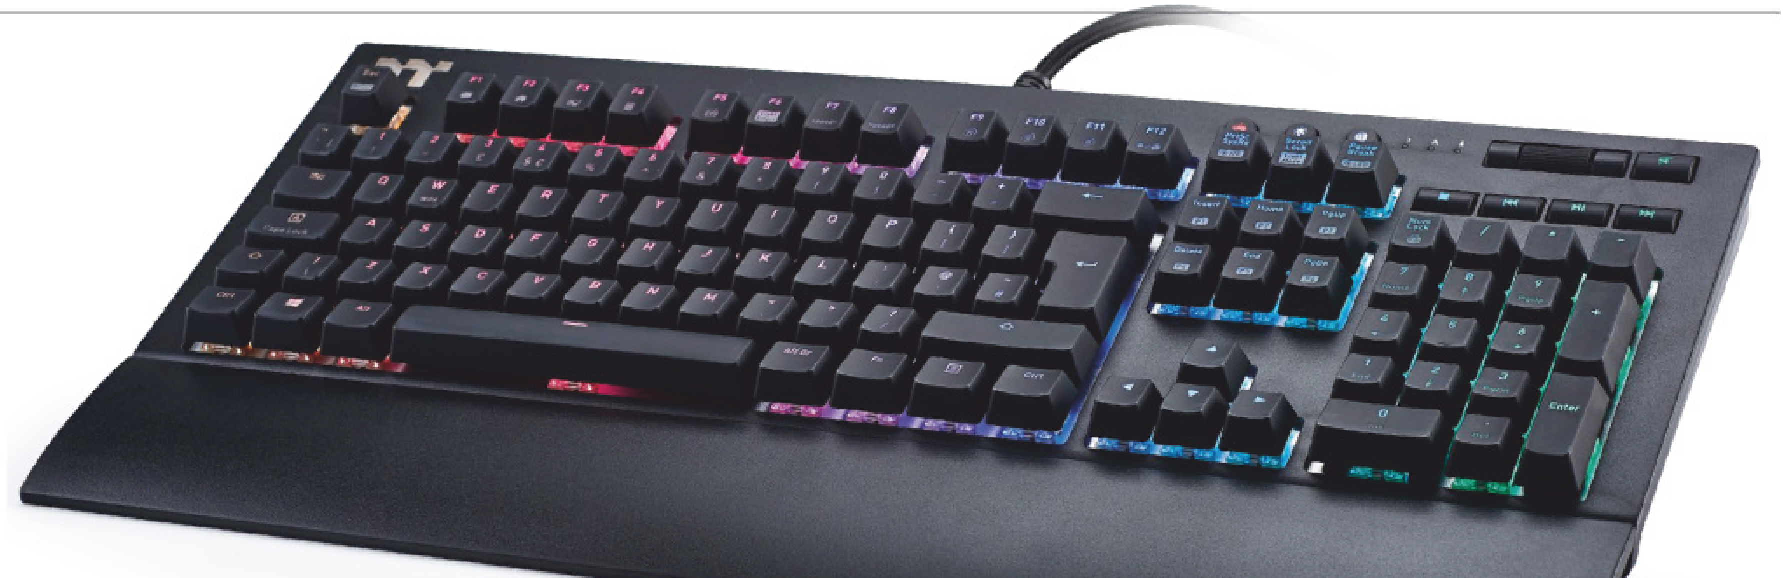

Thermaltake Tt Premium X1 RGB

The macro master with speedy keys

IF WE HAVE ESPORTS to thank for anything, it's that every manufacturer worth its salt (and with any money in the pot) is creating a division to dribble out gamer gear that, at least in some cases, pushes the envelope even for those who have never tasted Mountain Dew. Does the Thermaltake X1 RGB do anything particularly new? Well, yes, sort of. Does it need to? Absolutely not. But we're glad it exists.

The interesting bits, then. Top of the list, just because it's so silly, is the odd feature that enables you to control your keyboard's lighting and macros from your phone. It's a little bit paste-and-paper in its construction, and we're only 80 percent sure we trust its methods, given that it tunnels to your PC directly, but it's kind of neat to be able to adjust the twinklies with voice control, and even using Alexa. You can, though, pull off all of these functions directly on the keyboard using various combinations of Fn-whatever. This includes a super-neat macro programming function, messing with the speed and direction of built-in lighting effects, and even mixing colors to paint individual keys using F9–F12 as a palette. Basically, that cell phone software is frivolous at best and useless at worst, and we probably wouldn't even bother with the rather awful desktop app, given the huge level of control you have on the keyboard itself.

That lighting is pretty enough, packed with neat patterns, reflected off a silvery keyboard backing, and extending as far as the highlight logo on the top-left and the mute toggle top-right. There's a five-level brightness switch, which basically enables you to flip between dim, off, and back up to dim again—we appreciate not being blinded for once, if we're honest—and beside it are the macro-enabling game mode, and a Windows key toggle. A volume roller sneaks in above the dedicated media keys, always a nice touch, even when it's as rough and cheap feeling as this one, and in the box is a magnet-mounted plastic wrist rest, which is, we can exclusively reveal, definitely a plastic wrist rest. It didn't wander off anywhere while we were using it, so that's a plus.

Unusually, as in this is the first time this inveterate keyboard reviewer's got his fingers on them, Thermaltake has managed to bag Cherry's long-gestating MX Speed Silver switches for the X1 RGB. They actuate super-high, at almost half the distance of their closest cousins, the MX Reds, and they're similarly linear, meaning there's no bump, and little in the way of sound. Oddly, although they also match up perfectly with MX Red switches in terms of actuation force, Speed Silver switches feel much more pleasant, with the murky squidge of Reds replaced by a crisp, smooth, bouncy action. There are rough edges here, such as a wobbly backspace key, which sounds like it's made of aluminum foil, but it's an overall pleasing set of switches.

Checking more boxes, there's a pass-through for both USB and a four-pole headset jack, both of which are rare in other keyboards, so that's nice. Really, the X1 RGB does it all, and it's highly competent, so we're not sure whether to blame Thermaltake for the vague ennui it creates in us or not. Perhaps we're suspicious of those slight rough edges and janky software, perhaps we've seen so many great keyboards recently that a merely very good one just can't cut the mustard, or perhaps the X1 RGB is somehow not quite as exciting as the sum of its parts. More likely, we're just being grumpy. If you've been dying to get your hands on MX Speed Silvers, and you're one of those people who'll happily spend an hour poking around at your keyboard's lights, and a further two hours setting up macros for your favorite games, this is probably your dream keyboard—but it ain't cheap. **–ALEX COX**

VERDICT 8

Thermaltake Tt Premium X1 RGB

+ **SPEEDY** Pleasant switches; great on-board controls; pass-through.

− **NEEDY** Some rough edges; terrible software; slightly dim LEDs.

$140, www.thermaltake.com

SPECIFICATIONS

Switch Type	Cherry MX Speed Silver
Form Factor	Full-size
Media Keys	Yes
Macro Keys	Configurable
LEDs	Full per-key RGB
N-Key Rollover	Yes
Pass-Through	Four-pole 3.5mm, USB
Dimensions	18.3 x 6.7 x 1.5 inches
Warranty	Two years

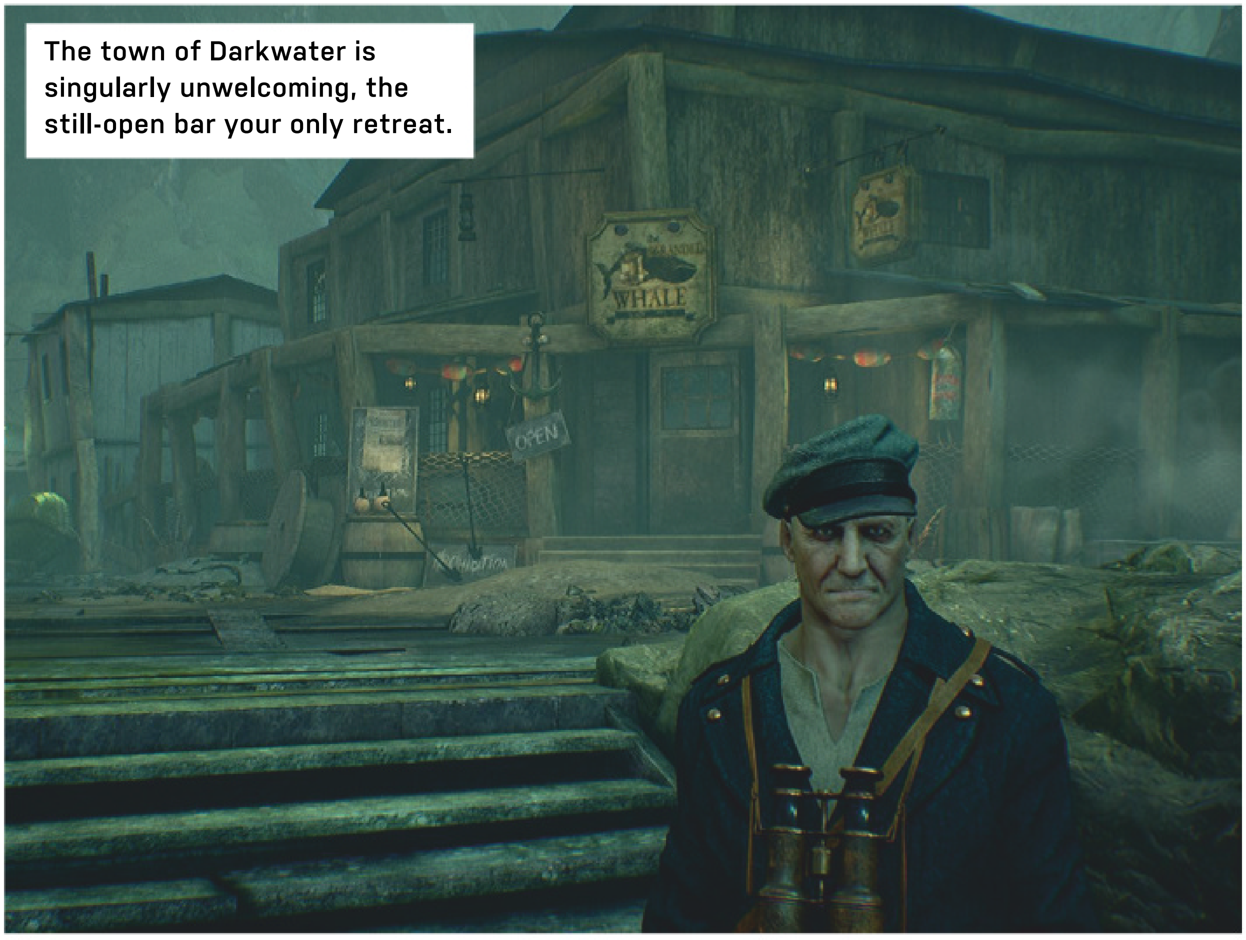

The town of Darkwater is singularly unwelcoming, the still-open bar your only retreat.

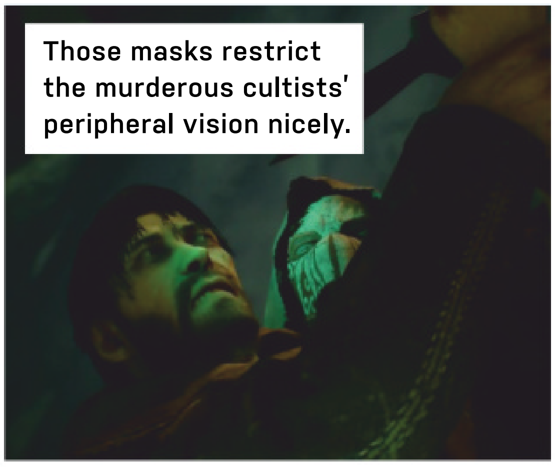

Those masks restrict the murderous cultists' peripheral vision nicely.

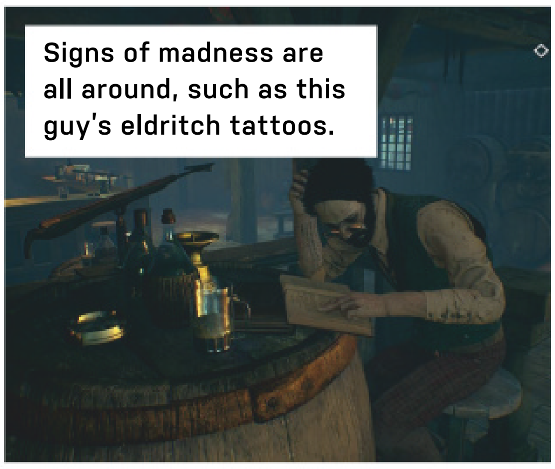

Signs of madness are all around, such as this guy's eldritch tattoos.

Call of Cthulhu

Ph'nglui mglw'nafh Cthulhu R'lyeh wgah'nagl fhtagn

H.P. LOVECRAFT may have done most of his writing almost 100 years ago, but he casts a long shadow over contemporary arts. While cosmic dread, non-Euclidean geometry, and sudden bouts of insanity don't always translate well to movies, they're perfect for games.

This particular stab at unknowable horror is based on the 1981 paper RPG of the same name, and sees grizzled private investigator Edward, suffering from PTSD following wartime experiences, look into a strange fire and disappearances on an island off the coast of Massachusetts. The era sees prohibition and the end of the whaling industry, so an isolated town full of men with too much time on their hands has found a new outlet for its baser urges.

As a study in foreboding atmosphere, *Call of Cthulhu* hits all the right unsettling notes. Faces are craggy, all deep-set eyes and filthy scars, and a strong maritime theme gives an excuse for the beams of a lighthouse to rhythmically pierce the misty gloom.

While exploring, a tale of a botched police investigation into the death of a family unravels, with Edward's private investigator skills digging up clues, and his medical knowledge used to interpret them. As you gain experience, points are added to your skills, but some can only be increased by reading books or finding artifacts. Most objects have a glowing white dot to guide you to them, but some are hidden, and if your "find hidden" skill isn't high enough, the game won't even tell you they're there.

Edward can also reconstruct past scenes, figures manifesting as ghostly recreations of often violent events that help his investigation and give him insight into the motives of the present. It's not the kind of horror game that tries to make you jump in the dark, more the sort that leaves you with a sense of unease, and a feeling that you're not in control.

Unfortunately, once the opening is over, it lets itself down. A wonderful scene in a mental hospital, which follows a tense close encounter of the tentacled kind, is followed by a mix of simplistic puzzles and enforced stealth as you escape. It even reuses puzzle mechanics, having you following cables or pipes to their destinations, while avoiding the kind of goons who could surely be punched out of the way. Get caught and you're dead. It recycles enemies, too, and there's only one section in which you engage in combat.

The feeling of helplessness in the face of unstoppable onslaught by creatures whose motives you can't understand might be Lovecraftian, but it's not great for videogames, where player agency is all. The paper RPG often left players dead or trapped in mental institutions, but here all you need to see the end credits is perseverance. Restarts are fast, and checkpoints sensibly placed, but sequences such as finding the right object in a room full of wrong ones before a monster gets you feel like trial and error, rather than tests of skill.

There have been several attempts to bring Lovecraft to the world of games, and this is one of the more successful. Leaving H.P.'s overwrought writing style behind in favor of damaged working class men and their struggle for significance, it's a horror game without the jump scares, and it brings a touch of the cosmic to small-town New England. **–IAN EVENDEN**

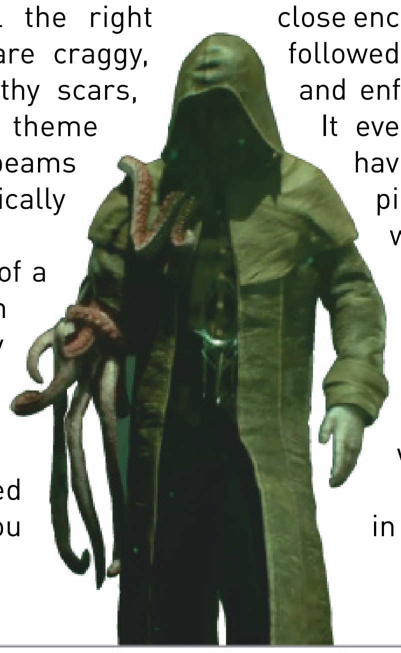

VERDICT 7

Call of Cthulhu

SHUB-NIGGURATH Relentlessly disturbing atmosphere; great tale of cults and monsters; some really craggy faces.

NYARLATHOTEP Enforced stealth; one-hit kills; trial and error.

RECOMMENDED SPECS Core i7-3820/FX-8370; 8GB RAM; GTX 970/R9 390.

$40, www.callofcthulhu-game.com, ESRB: M

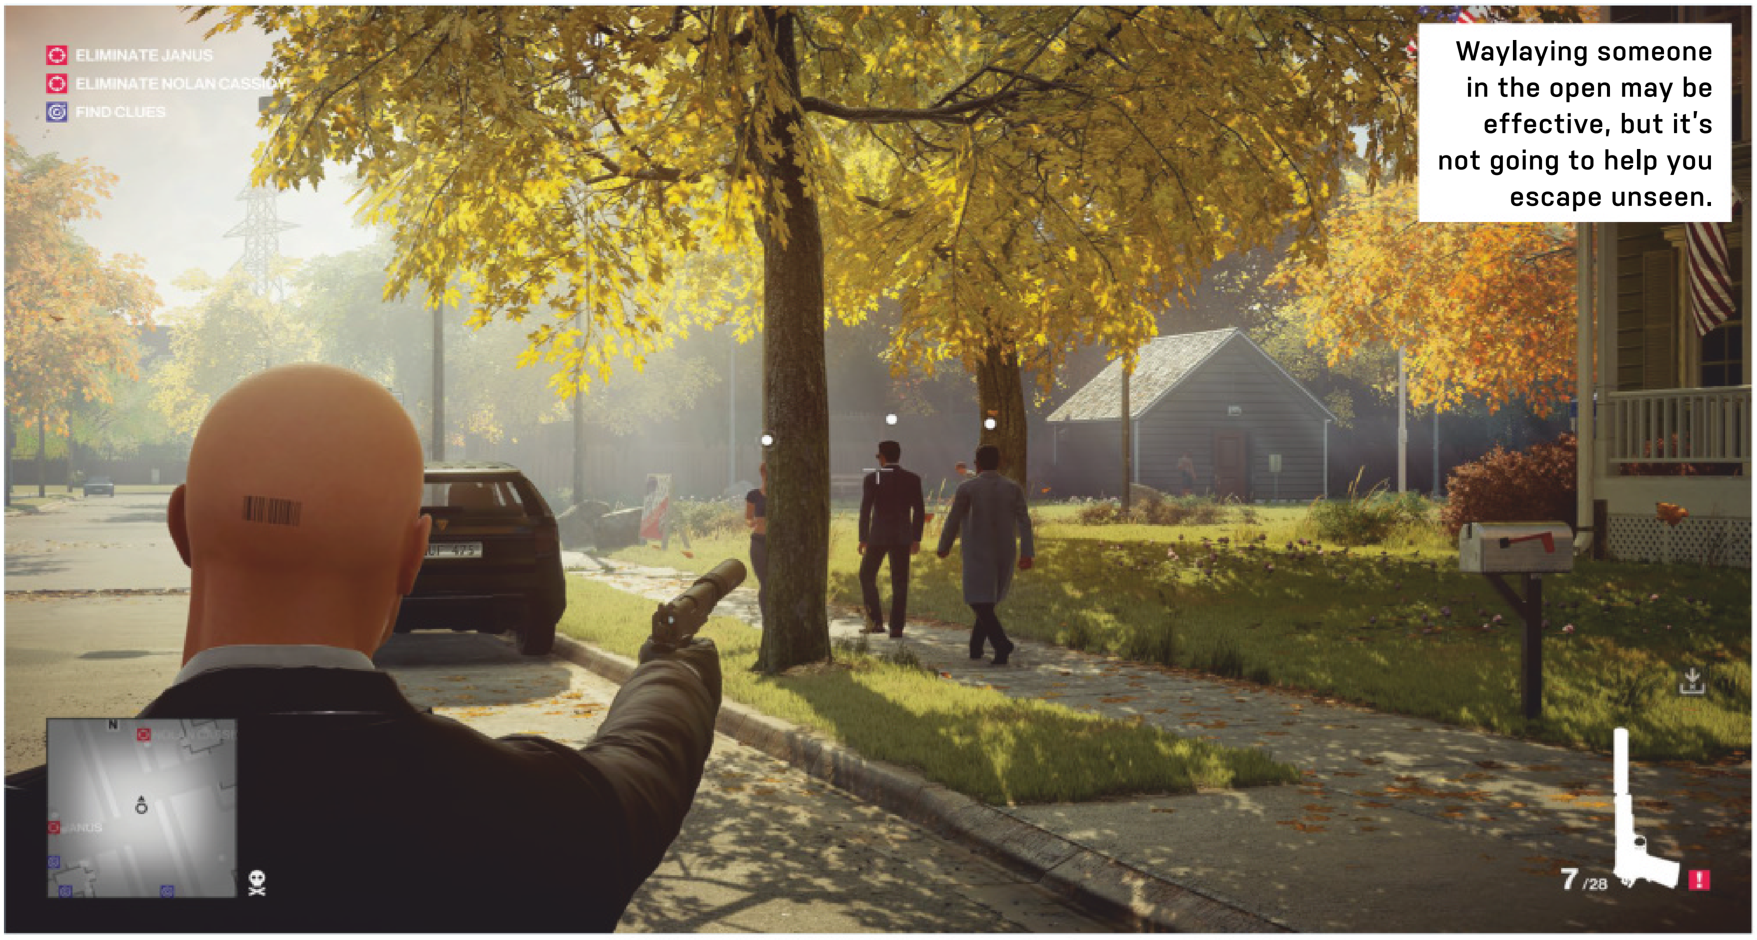

Waylaying someone in the open may be effective, but it's not going to help you escape unseen.

Hitman 2

Man compensates for lack of hair with weaponry

A HITMAN LEVEL isn't designed to be played just once. There's the initial explore, when you go in knowing just the intel you've been supplied with beforehand, and need to scope everything out. The locations of weapons, targets, and security cameras are still unknown, and it's quite possible that you'll need to restart before your target is cold.

Then there are the practice runs, when you know the layout, and are experimenting with murder methodology, graduating from simple suppressed gunshots to poison in drinks to elaborate creations involving gas pipes, overloaded valves, and fireworks. Later, you're playing it on a timer, trying to shave seconds off to knock your rival from the top of the leaderboard. By this point, you know the level inside out, counting under your breath as you wait around a corner for the exact number of heartbeats it takes for a guard to turn his back.

This is what sets *Hitman* apart from other games: It's meant to be repeated, to be perfected. To thrash through a level and escape with a simple kill, and maybe some surveillance video left behind, isn't good enough. You need to really get inside the level designer's mind, to see what they intended, and pull it off with a flourish. *Hitman* isn't a third-person shooter or even really a stealth game, it's a Rube Goldberg machine, waiting for you to manipulate it. And it's for this reason that the previous game's release as a series of episodes made sense. Without the pull of the next level enticing you to speed up, you could spend a month patiently laying plans, and watching them pay off, then refining them on another run.

Hitman 2 gives you all its levels in one go, and you can even play them out of order. The new scenarios feel more ambitious, something that can be said of almost every game in the series, which has always tried to outdo its own past.

KILL JOYS

There's a vague plot involving international terrorists, but *Hitman 2's* levels are so full of characters and stories, disguises and machinations, that they feel self-contained, more like episodes than the constituent levels of a whole. There's suburbia, full of identical houses, but owned by different people with stories to tell, and these come out as you wander through. Then there's a motor race, with two targets and huge crowds, but also with a warren of back corridors and staff-only areas populated by people doing their jobs, and sometimes doing nasty things to one another. The crowded streets of Mumbai offer three targets, but also a street gang, barber, tailor, laundry, and diverse mundanity that can be repurposed for killing. Elsewhere, you'll come across Colombian drug lords and an island paradise for the rich, underlining the series's preference for killing the wealthy whenever possible. There are even some robots to fiddle with.

This is iteration rather than revolution, something further evidenced by the ability to play 2016 *Hitman* levels within the game. It's like a GOTY edition of that game, one that comes with all DLC. It's a novel approach, and one that will supply anyone who's not played the previous game with a bonanza of content. A lot of missions feel like riffs on those from previous games, so perhaps the devs are refining their approach to the game, too. One thing that's more than a riff is the tutorial, which is identical to 2016's, firming up the view that this is season two, delivered all at once for binge watchers. Or mass murderers. **–IAN EVENDEN**

VERDICT

8

Hitman 2

+ CREED Refined iteration of the previous game with new levels; complex and confident.

− ROCKY Not too much in the way of new ideas; occasional glitches in NPC routines.

RECOMMENDED SPECS CPU i7-4790 or equivalent; 16GB RAM; GTX 1070 or RX Vega 56.

$60, www.hitman.com, ESRB: M

LAB NOTES

JEREMY LAIRD, CONTRIBUTING EDITOR

Tweaking Our Take on SSDs

It's time to adjust our aim at an ever-moving target

SSDs ARE A PAIN in the ass. At least, they are from a lab-testing perspective. Of all the major components, they're the tricksiest to test. For instance, it's simply not viable to assess long-term performance of an SSD when you might only be loaned the darned thing for a few days.

Likewise, the metrics by which SSDs are measured have shifted dramatically over time. At first, it was all about raw bandwidth. Then it became apparent that drive performance could degrade over time. The latter issue has been largely resolved, with well-engineered drives at any rate. More recently, it's become clear that what really matters for desktop PCs is 4K random access at low queue depths.

You can read more about that in our feature on page 36, but suffice to say it's time to tweak our testing methodology. We were already measuring 4K performance, but now we're increasing granularity by testing 4K reads and writes at queue depths of one, two, and four. We're keeping sequential reads and writes to provide a headline view of peak performance.

Likewise, our 30GB internal file copy test remains a useful real-world test. Admittedly, any half-decent NVMe drive will chew that benchmark up, but cheaper SATA drives that lack DDR cache or rely on a small chunk of SLC cache to post decent-looking sequential numbers can be caught out. An example of that is Crucial's new entry-level BX500 drive, particularly in 120GB.

That drive can only sustain its advertised performance for about 1.5GB of sequential data, at which point throughput falls off a cliff. Taken together, our tests can't quite provide that full long-term view of how a drive will stand up to years of use, but the mix of sequential and random synthetics, plus a real-world sense check in the form of our internal file copy test, give a good overall picture of what you can expect.

intel
Intel® Optane™ SSD 905P

Cutting-edge drives demand new metrics.

ALAN DEXTER

Executive Editor

Looking at the statistics for Black Friday and Cyber Monday highlighted a couple of things. Firstly, that lots of us are searching for bargains on the very latest gear—no real surprise there—and such bargains are incredibly rare. The other big takeaway was that there were some incredible deals on the storage front: SSDs (both SATA and NVMe) are much more affordable than they were a year ago anyway, and with a discount on top of that, plenty of us could be seen reaching for our credit cards to give our systems a huge boost to performance and capacity. The really good news is that it looks like these price tumbles are going to continue through 2019.

JOANNA NELIUS

Hardware Staff Writer

I've recently reviewed a few pre-built systems—Asus ROG Strix, Corsair Vengeance, Walmart OverPowered—but the most interesting one was an NZXT system I designed myself. Call me old, but back in my day, there were no fancy custom pre-built services. During the pre-Internet days, we'd go the store and pick out a computer. Now there are companies that build a PC to your specifications, then ship it to you. This is a new discovery for me, because I only recently started writing about hardware. As much satisfaction as I get from building my own PC, my time grows increasingly limited, so I might give NZXT or another such service a try when I build my next PC.

Editors' Picks:
Digital Discoveries

Jarred Walton, senior editor, and Bo Moore, hardware lead, enjoy the wider world of technology in their homes

INSTANT POT ULTRA 8-QUART

Little known fact: I'm a family man, and I can cook. (Don't tell my wife.) I've made plenty of meals in crock pots over the years, but there's one small problem: time. If I remember to get things ready in advance, letting things simmer on low all day is great. Often I forget until a few hours before it's time to eat. Did anyone want pizza for dinner instead?

But the Instant Pot brings pressure cooking to the mainstream, and in doing so, cuts cooking times way down. Take one of my favorites: roast beef with pepperoncini and some seasonings. With a slow cooker, I'd cook it on low for 10–12 hours, or in a pinch I could try to get it ready to eat in 5-6 hours on high. But cooking on high also tends to make the meat less tender, which runs counter to the slow cooker mantra. In the Instant Pot, thanks to pressure cooking, it takes about 75 minutes.

We got the 8-quart one; it's awesome. My wife loves it, and I'm wondering why I didn't give in to her request earlier. Probably because I'm a cheap knucklehead. Do yourself a favor and buy an Instant Pot—you can get rid of your crock pot and rice cooker while you're at it.
$180, https://instantpot.com

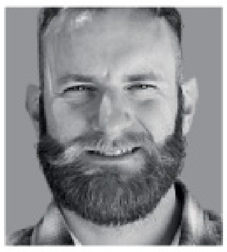

PHILIPS HUE WHITE AND COLOR TWO-BULB STARTER KIT

This Black Friday weekend, I finally took the plunge and invited Jeff Bezos into my home. Which is to say, I picked up an Amazon Alexa device. To go along with it, I grabbed a Philips Hue two-bulb starter kit, so now I can control my two living room lights like a wizard. The starter kit comes with two bulbs (or four in the more expensive kit), plus the Hue bridge needed to control it all. With that single bridge, I can add additional bulbs or other lights, like lightstrips or spotlights, and control them all with Alexa.

Doubly cool, Philips Hue can integrate with Razer Synapse, which I begrudgingly use—despite all the annoying updates—because a Razer mouse and keyboard are my daily drivers. The software can be frustrating, but it does a fine enough job controlling the lights on my mouse and keyboard. With Hue integration, Synapse can control my living room lights as well, synching them up to the color spectrum cycle of my other peripherals. I've even set up a custom command to switch my lights from normal to a jubilant blast of colors. Who cares if Lord Bezos might be spying on me 24/7? I'm a rainbow wizard, it's great.
$60, www.meethue.com

GEEK TESTED & APPROVED

Samsung Wireless Fast Charging Stand

WHEN IT COMES to smartphones, I swear by the wireless Qi connection standard. Having been introduced to it with the Samsung Galaxy S7 Edge, and continuing to use it with the S9 Plus, without a shadow of a doubt, I couldn't live without it now.

I litter a plethora of these bad boys all over the house and office. I have a flat one on my desk at home, an angled one at work, another angled one by my bedside table, and even one downstairs next to the TV for when I'm binging on TV series with my future wife.

The thing is, they're so ridiculously convenient—there's no chance of damaging the port, a reduced chance of damaging the cable, and keeping my phone charged, quickly and easily, is an absolute dream. I get up, go to work, I'm on 100 percent charge. Leave the office, still on 100 percent. Jump in the car, 100 percent. You get the picture. Battery life doesn't affect me. On top of that, because it all functions off the Qi wireless standard, they're cross-compatible, too. From iPhones to LG to Samsung—take your pick.

There are only two downsides I can think of, the first being the price. As standard, these come in at $70, although more often they're $30–$50 when on sale. The second is the fan. As the quick chargers generate a bigger magnetic sphere, they create more heat, and require a tiny fan to dissipate it, which can be annoying if you're a light sleeper. **–ZS**
$54, www.samsung.com

LETTERS

WE TACKLE TOUGH READER QUESTIONS ON...

> Smooth Rewind
> Dual Keyboards
> Screen Deep Dive

Visible Search

I've read lots of discussions and reviews of various streaming devices and the apps in TVs and DVD/Blu-ray players, and they never seem to discuss a few very basic but, to me, extremely important usability issues, such as visible search, frame-stepping, and other related "transport" issues. I've always been horrified whenever I've tried to watch streamed content—how can anyone search without being able to see if the video you're fast-forwarding or rewinding through is OK? I have a new Sony UltraHD Blu-ray and tried the Amazon Prime app, and no surprise—no visible search. If I try to replay the last few seconds because I missed a word or just want to watch a cool bit over and over, I have to check the time as it goes into rewind, which it does at about 5-6 seconds per second. If you don't, it's easy to get lost, and trying to get back to where you were can be a nightmare. If you hit rewind twice, it moves at something like two minutes per second!

The icing on the cake is that after you manage to hit "Play" at the right time, say five seconds from where you were, it takes seven to eight seconds before it manages to start playing again. Who thinks that's even close to usable? You had to have the cheapest crappiest VHS VCR to not have visible search. Do any of the standalone streaming devices do this any better? My experience is limited, and I don't have the energy or the sanity reserves to keep trying to find one that meets the most basic requirements for me. I can't stand to watch good movies or television with such limitations. Daum PotPlayer gives you four levels of one-key stepping, forward or backward, you can grab the cursor and drag to whatever time you want, and see the picture you're dragging through. Even YouTube in a browser is far superior. Am I missing something? How can anyone put up with such horrors? **–Rob S**

EXECUTIVE EDITOR, ALAN DEXTER, RESPONDS:
I agree that some of the user experiences with streaming can be sub-optimal, but the convenience of being able to browse so much content from the comfort of your couch often outweighs them. I use a big smart 4K Samsung TV as my main screen at home, and don't mind the searching, but maybe I'm just lucky. Still, as you've pointed out, PCs are far better at such things, so if frame-by-frame searching is your thing, a media center PC is the best solution.

Split Keyboards

Perhaps you could forward this to Alex Cox, author of the "Create Your Own Macro Keyboard" article in the September issue. It's a complex process, which I might pull off. I'd like to know whether one can disable its settings or not. Would I be able to toggle the interception on/off? **–M. Lutz**

CONTRIBUTING WRITER, ALEX COX, RESPONDS:
There's a number of ways to disable macros when using AutoHotKey and the Interceptor method we described—primarily, you could just stop the AHK script from running. Your macro keyboard will still send Key-F23, but as there's almost no Windows software, bar AHK itself, that listens for it, it will behave like a normal keyboard. Killing the intercept.exe process has the same effect, but cuts out the F23 portion—run it again, and it loads your last configuration.

Open Correction

In the December 2018 issue, the article on open-source processors has a comment called "An Open Dream." I thought that comment missed the mark on a few points. First, Arduino is not a microcontroller range, it is a microcontroller platform. The difference may seem insignificant but it is actually quite large. Arduino does not make microcontrollers. Arduino used existing microcontrollers made by Atmel and placed them on a board. Then it created a coding and programming environment to form a platform for development. Atmel (now owned by Microchip) does not make an open microcontroller. Second, while Arduino's

↘ submit your questions to: comments@maximumpc.com open nature about its board designs is great, it is not the main thing that contributes to its popularity. When using its platform, Arduino lowered the barrier of entry to coding for microcontrollers. It hid away all the yucky hardware stuff behind many libraries. This library approach continued as other people built more and more libraries. There are hundreds if not thousands of libraries and examples to simplify coding. Finally, the image shown above the note is not even an Arduino board. The image is of what the Arduino platform calls a shield. A shield is an expansion board that plugs directly into the Arduino board to give it extra functionality, which is then further covered by libraries. This particular image is of an Ethernet shield. **–Igor Vinograd**

EXECUTIVE EDITOR, ALAN DEXTER, RESPONDS: Thanks for your clarification. And yes, we may have been a little off with our image selection, although the idea was simply to show an Arduino device. The essence of open source providing another possible future for processors still stands, though. And the key idea you've pointed out, that a lot of the hardware managing is done for you through libraries, is vital to the success of any future processor, or indeed open-source hardware. We're still excited by the possibilities.

Screen Scene

I have been reading *Maximum PC* since 1999. It has shaped my love for PCs and helped me with my builds. With the rash of new video card tech, there is a debate about which monitors to purchase. There are so many options.

Could you guys dedicate an issue and discuss and explain about all the tech behind today's monitors? Such as what the differences are between panel types (VA, IPS, TN, etc), HDR, OLED, LCD, and LED. Talk about response rates (GtG), contrast, pixel density, native resolution, G-Sync, FreeSync, frame tearing, 1080p, 4K, bezel, aspect ratios, backlight, pixel pitch, viewing angles, display colors, VESA standards, flat versus curved, connectivity types (USB C, HDMI, DisplayPort), and anything else I might have missed.

It's quite overwhelming, and I would love to see you guys take on the challenge, and explain all that jargon to readers like me, who are trying to make sense of it all. **–Chad Schoenthaler**

EXECUTIVE EDITOR, ALAN DEXTER, RESPONDS: We did take a look at the current state of the monitor market last issue, and we covered most of the things you mention here, along with recommendations for screens at certain budgets. However, if you want a deeper dive into screen tech, let us know, and we'll see what we can do.

Buying Advice

Please tell me where Alan Dexter purchases his components (ref: "Tracing the Rays of the Future," page 70, December 2018 issue). I would like to build the machine detailed in his article. **–Jack Williams**

EXECUTIVE EDITOR, ALAN DEXTER, RESPONDS: We use a variety of sources for piecing together the builds each month, although for the main part, we get our gear from manufacturers sending us the hardware directly. Even so, if you use a site like PC Part Picker, you can see where the best deals are, and what you'd expect to pay for such components (prices are pulled in from several sites, including Amazon, Newegg, and the like).

For this system, you can see what we've pieced together by pointing your browser at https://pcpartpicker.com/list/zC6HsZ, although there are a few things to note, thanks to plenty of the items being out of stock pretty much everywhere at the moment. The Raijintek Orcus 240 is too new to appear on this list as well, although by the time we go to print, hopefully this will be rectified. You can replace individual parts easily enough, and search for better deals as you do so.

Budget Workstations

I've subscribed for two years and love the articles, especially the print format, which no other PC mag offers anymore. The article some time back about buying used PC equipment was excellent. I replaced a Core 2 Quad Dell with a used HP Z420 from eBay that had one of Intel's "red-headed step children" in it—a Xeon E5 2670 (eight cores with Hyper-Threading)—plus 32GB of RAM, and a full-size case. All I had to add was my Quadro K620. It was far cheaper (under $400) than an Intel Core i7 system, and has plenty of room to expand. Why not run an article about Intel's Xeon chips, and while you're at it, another article on Nvidia's Quadro processors? Aside from some brief mentions, I've seen nothing on these "step children" in *Maximum PC*. I'm sure there are readers who use them. Mine get used daily running DAZ Studio Pro 4.10. That's another thing you should cover—it's a 3D rendering program similar to Poser, but (for now) the Pro version is free. It even has a ray-tracing rendering engine called Iray that produces incredibly life-like results. They're only single images, but you can't beat the price. **–Eddie Davis**

EXECUTIVE EDITOR, ALAN DEXTER, RESPONDS: Good points well made, and definitely something we'll look at in the future. There are some great deals to be had out there if you know what you're looking for, and it certainly pays to know what your software will support before reaching for your wallet.

Build your own ray-tracing monster—if you can find the components.
GEFORCE RTX

MAXIMUM PC DIGITAL EDITION
ANYTIME, ANYWHERE!

AS LITTLE AS $1.99 PER ISSUE!
MAXIMUM PC
2019 TECH PREVIEW
Next-gen hardware from AMD, Intel, & Nvidia that will transform your PC
PLUS! Build a 1080p gaming rig
THE BEST GAMES COMING THIS YEAR

Delivery of **MAXIMUM PC**, in your choice of
Apple Newsstand, Google Play, or Zinio

From only $1.99 a month

Search for **MAXIMUM PC**
on your device's app store to download now

Available on the App Store GET THE MAGAZINE ON Google play zinio™

Sponsored by newegg.com
TAKE IT FROM A GEEK.℠

THE BUILDS

BUDGET

INGREDIENTS

PART		PRICE
Case	BitFenix Nova TG	$45
PSU	450W EVGA 450 BT 80+ Bronze **NEW**	$39
Mobo	ASRock AB320M-DGS	$50
CPU	AMD Ryzen 3 1200	$90
GPU	Gigabyte GeForce GTX 1050 Ti 4GB **NEW**	$170
RAM	8GB (2x 4GB) Corsair Vengeance LPX @ 3,000MT/s **NEW**	$76
SSD	120GB SanDisk SSD Plus **NEW**	$24
HDD	1TB Hitachi GST Ultrastar A7K1000 **NEW**	$35
OS	Ubuntu Desktop Linux 18.04 LTS 64-bit	$16

Approximate Price: $545

IT'S BEEN A LONG TIME since we've changed anything dramatic in our budget build, and for good reason: There's nothing right now that would shake up its performance. Everything seems to be situated at the high end. Even Nvidia's RTX 2060 is still waiting in the docks for some meaningful competition, trying to eke out that last cent for folks committed to buying the latest-gen GPUs. That said, we've had a bit of a shuffle around, saved some cash, and managed to splurge on a slightly better GPU: the GTX 1050 Ti. It's not frame-breaker, but it's a darn sight better than the 1050. We're talking a 20 percent improvement on the base hardware alone. Unfortunately, it's not the 80 percent you'd get for going to the 1060 3GB model, but still an improvement nonetheless. This was possible thanks to memory, SSD, HDD, and PSU price drops, and amazingly this build is still $10 cheaper than last issue,.

MID-RANGE

INGREDIENTS

PART		PRICE
Case	Fractal Design Meshify C	$90
PSU	550W EVGA SuperNOVA 550 G3 80+ Gold	$70
Mobo	Asus ROG Strix B450-F Gaming	$110
CPU	AMD Ryzen 7 2700	$280
Cooler	Corsair Hydro H100i Pro 240mm RGB AIO	$100
GPU	Zotac GeForce GTX 1070 Ti AMP! Edition **NEW**	$360
RAM	16GB (2x 8GB) G.Skill Aegis @ 3,000MT/s	$100
SSD	250GB Samsung 970 Evo M.2 PCIe SSD	$78
HDD	1TB Hitachi GST Ultrastar A7K1000 **NEW**	$35
OS	Windows 10 Home 64-bit OEM	$100

Approximate Price: $1,323

WE DESPERATELY WANTED to trade out that 1070 Ti for a 1080, but the price difference is nearly double that of the Zotac card we found. "Frustrating" doesn't quite capture the feeling. "Confusing" might be a better word. In our testing, we've found that the 1070 Ti is only around three to five frames per second slower than the 1080. All in all, though, this issue saw little change, apart from that new GPU and the HDD. The lingering tendrils of Black Friday still seem to be creeping into our mid-range rig, and we've managed to save a grand total of $80 in the process. Not too shabby, huh? Next on the upgrade agenda would be a better-specced motherboard, potentially the Asus Prime X470-Pro. Because it comes in at $177 right now, that's a $67 increase, which would still keep us below our $1,400 budget, but also provide some seriously improved connectivity.

TURBO

UPGRADE OF THE MONTH

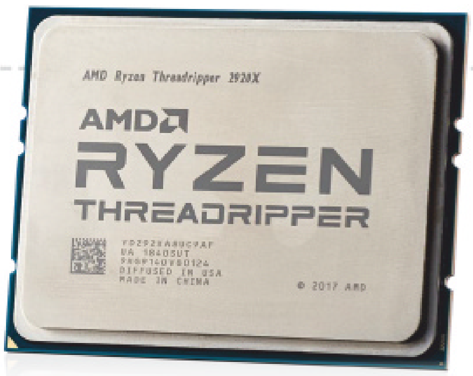

AMD THREADRIPPER 2920X

We're seriously impressed by AMD's Threadripper lineup, having just about wrapped up our second-generation TR testing. All in all, it's a sound investment for anyone looking to get into video editing, or any similar professional pursuit. The 2920X, in particular, has caught our attention, not only for its price point, but also because of its versatility. Plus, when you look at how far AMD has come in such a short space of time with memory improvements and the plug-and-play nature of its Threadripper lineup, it's hard not to feel admiration. If we were looking to invest in a new rig, the TR 2920X would be our first port of call. With its multicore performance, competitive in-game prowess, incredible temps, and reasonable price, it's an awesome chip. And when paired with a decent cooler, some good memory, and a solid motherboard, it becomes a true monster.

$650, www.amd.com

AGAIN, NOT A HUGE BATCH of changes in our Turbo build. Annoyingly, despite the fact we suspect Nvidia and AIBs have excessive levels of GTX 1080 Ti stock, the prices for the cards are sky-high compared to the RTX 2080 right now. Given the two cards are the same, it seems Nvidia is attempting to push the market into its ray-traced field, likely to bolster the attractiveness of its ray-tracing gameworks features to developers, and in turn get more sales of the RTX 2080 Ti, its true money-maker. All speculation, but it seems likely.

Outside of our GPU woes, everything else remains relatively the same. You could save some serious cash by dropping out the 2950X for a 2920X. You'd lose out on four cores, but in game, at least, you'd still be packing some seriously impressive frame rates, and 12 cores aren't to be sniffed at—that's still 24 threads.

We're finally starting to see some serious price drops in the memory department. Good quality kits are falling by $20–$30 at a time, and it's looking good for the foreseeable future, with China threatening class action lawsuits against price fixing. Aside from that, we've kept it all fairly standard. We have heard tell of our favorite CPU cooler having some faulty review units from Newegg, so we'd keep an eye on that (perhaps opt for a Corsair or NZXT variant instead, if you are worried), but otherwise we're happy with where this build's at.

For more of our component recommendations, visit www.maximumpc.com/best-of-the-best

INGREDIENTS

PART		PRICE
Case	NZXT H700i	$172
PSU	Corsair RM750x White 750W Gold Certified	$100
Mobo	Asus Prime X399-A	$300
CPU	AMD Threadripper 2950X	$900
Cooler	Enermax Liqtech TR4 II 240mm RGB AOI	$133
GPU	Gigabyte GeForce RTX 2080 Gaming OC	$750
RAM	32GB (4x 8GB) Corsair Vengeance LED DDR4 @ 3,000MT/s **NEW**	$270
SSD	512GB Samsung 970 Pro M.2 NVMe	$178
HDD	2x 2TB Seagate BarraCuda ST2000DM006	$120
OS	Windows 10 Home 64-bit OEM	$100

Approximate Price: $3,023

Maximum PC (ISSN 1522-4279) is published 13 times a year, monthly plus a Spring issue following the March issue, by Future US, Inc., 11 West 42nd Street, 15th Floor, New York, NY 10036, USA. Website: www.futureus.com. Periodicals postage paid in San Bruno, CA, and at additional mailing offices. Newsstand distribution is handled by Curtis Circulation Company. Basic subscription rates: one year (13 issues) US: $24; Canada: US$40; Foreign: US$40. Canadian and foreign orders must be prepaid. Canadian price includes postage and GST (GST #R128220688). PMA #40612608. Subscriptions do not include newsstand specials. POSTMASTER: Send changes of address to *Maximum PC*, PO Box 5852, Harlan, IA 51593-1352. Future US, Inc. also publishes *Mac|Life*, and *PC Gamer*. Entire contents copyright 2018, Future US, Inc. All rights reserved. Reproduction in whole or in part is prohibited. Future US, Inc. is not affiliated with the companies or products covered in *Maximum PC*. Reproduction on the Internet of the articles and pictures in this magazine is illegal without the prior written consent of *Maximum PC*. Products named in the pages of *Maximum PC* are trademarks of their respective companies. PRODUCED IN THE UNITED STATES OF AMERICA.

22 THOUSAND

TENNIS COURTS EVERY DAY!

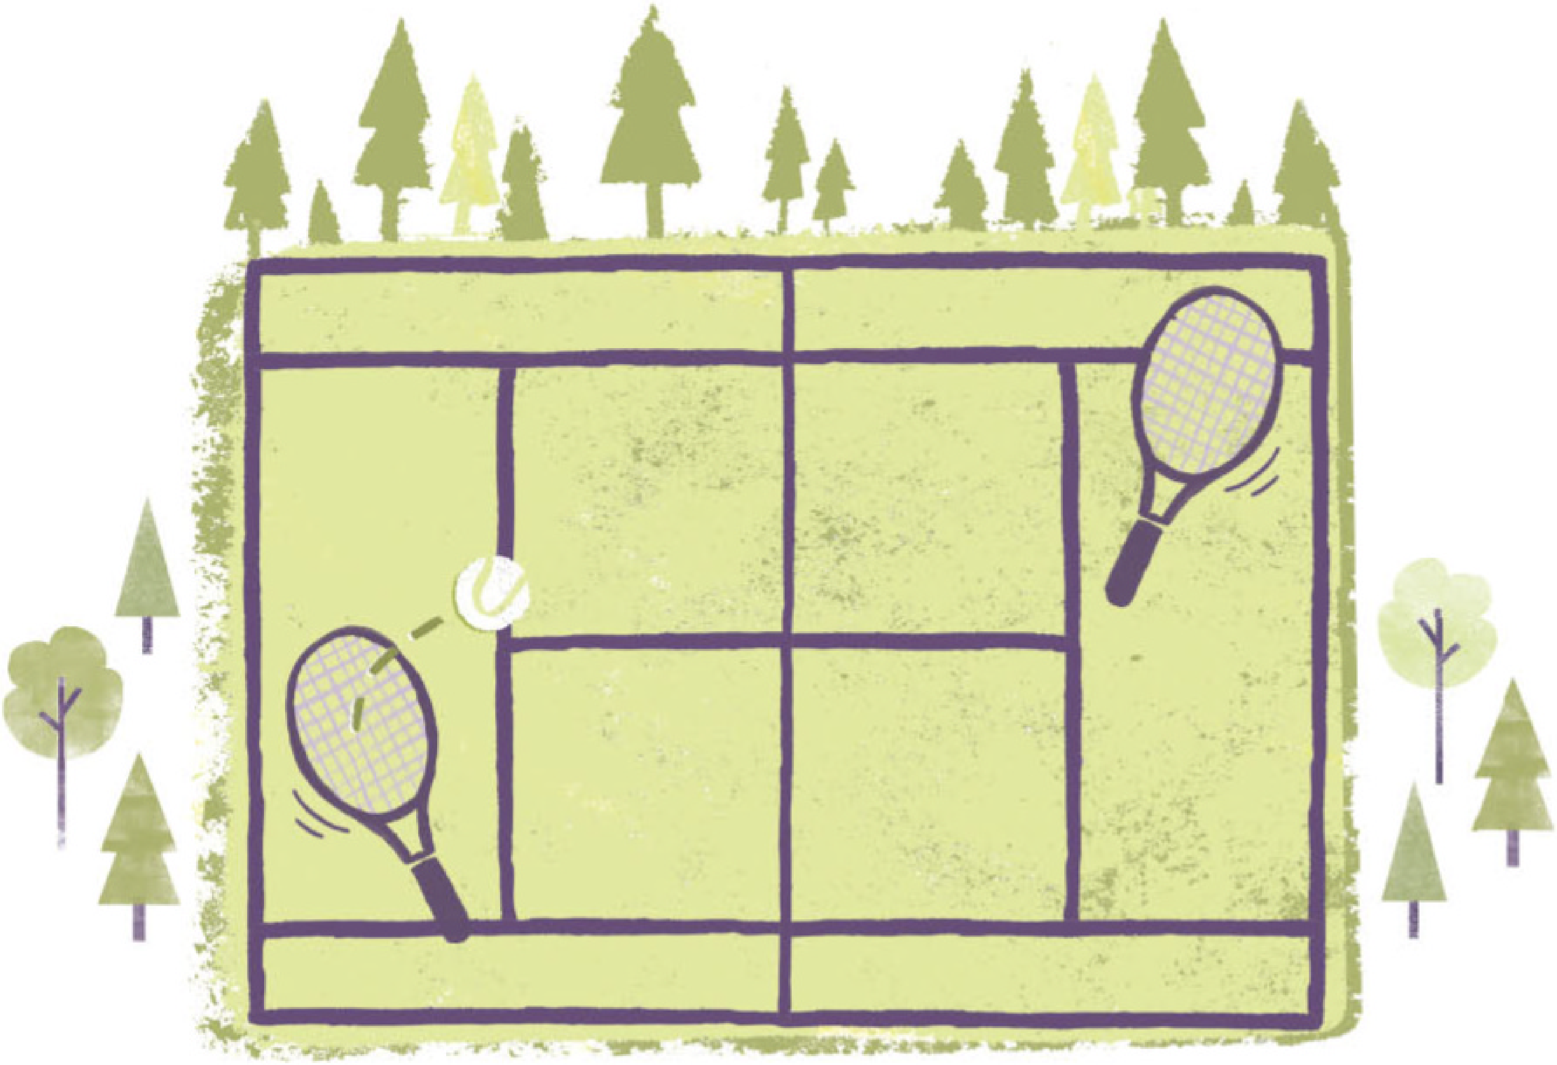

Did you know that U.S. forests, which provide wood for making paper and many other products, have been growing every day, equal in area to more than 22,000 tennis courts!

Love magazines? You'll love them even more knowing they're made from natural, renewable and recyclable wood.

UNFAO, Global Forest Resources Assessment 2005-2015

Two Sides is a global initiative promoting the responsible use of print and paper which, when sourced from certified or sustainably managed forests, is a uniquely powerful and natural communications medium.

Discover
more great
#PaperFacts at
twosidesna.org

TWO SIDES

PC GAMER
EDITOR'S CHOICE
MAXIMUM PC
KICK ASS!

We build the world's most advanced PCs.

Experience a new level of performance with an award-winning Digital Storm PC. Built with the latest technology, highest quality components and backed by lifetime support. Visit our website and build your dream PC today.

DIGITALSTORM

LEARN MORE: WWW.DIGITALSTORM.COM

Digital Storm PCs featuring Intel® Core™ i7 processors.

intel CORE i7 inside

The Digital Storm Logo, and 'World's Most Advanced PCs' are trademarks of Digital Storm. Intel, the Intel Logo, Intel Inside, Intel Core, and Core Inside are trademarks of Intel Corporation in the U.S. and/or other countries.